OAK RIDGE NATIONAL LABORATORY

MANAGED BY UT-BATTELLE FOR THE DEPARTMENT OF ENERGY

ORNL/TM-2010/149

U.S. Department of Energy, Office of Science

High Performance Computing Facility Operational Assessment, FY 2010 Oak Ridge Leadership Computing Facility

August 2010

Prepared by

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

UT-BATTELLE

ORNL-27 (4-00)

ORNL/TM-2010/149

U.S. Department of Energy, Office of Science

HIGH PERFORMANCE COMPUTING FACILITY OPERATIONAL ASSESSMENT, FY 2010 OAK RIDGE LEADERSHIP COMPUTING FACILITY

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

August 2010

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

CONTENTS

Page

LIST OF FIGURES

Figure **Page**

LIST OF TABLES

Table **Page**

ACRONYMS

3-D	three-dimensional
ACM	Association for Computing Machinery
ADIOS	ADaptable Input/Output System
ALCC	ASCR Leadership Computing Challenge (DOE)
ALCF	Argonne Leadership Computing Facility
ANL	Argonne National Laboratory
API	application programming interface
ARC	Appalachian Regional Commission
ARRA	American Recovery and Reinvestment Act (of 2009)
ASCAC	Advanced Scientific Computing Advisory Committee (DOE SC)
ASCR	Advanced Scientific Computing Research (DOE program office)
C&A	certification and accreditation
CAM	Community Atmosphere Model
CCSM	Community Climate System Model
CFD	computational fluid dynamics
CHiMES	Coupled High-Resolution Modeling of the Earth System
CLE	Cray Linux Environment (operating system)
COV	Committee of Visitors
CR	continuing resolution
CSB	Computational Sciences Building (ORNL)
CSPP	Cyber Security Program Plan
CUDA	Compute Unified Device Architecture (parallel computing architecture)
CY	calendar year
DD	Director’s Discretionary
DDT	Distributed Debugging Tool (Allinea Software Ltd.)
DDN	DataDirect Networks (data storage infrastructure company)
DME	development, modernization, and enhancement
DOE	Department of Energy
EOY	end of year
eSimMon	electronic Simulation Monitoring
ESMF	Earth System Modeling Framework
FAQ	frequently asked question
FIESTA	Framework for Integrated End-to-End Scientific Data Management Technologies for Applications
FTE	full-time equivalent
FWP	Field Work Proposal
FY	fiscal year
GB	gigabyte
GB/s	GB per second
GFDL	Geophysical Fluid Dynamics Laboratory

GPU	graphics processing unit
GROMACS	GROningen MAchine for Chemical Simulations
GTC	Gyrokinetic Toroidal Code (fusion code developed at Princeton Plasma Physics Laboratory)
HBCUs	Historically Black Colleges or Universities
HPC	high-performance computing
HPCF	high-performance computing facility
HPCFOA	high-performance computing facility operational assessment
HPSS	High-Performance Storage System
I/O	input/output
INCITE	Innovative and Novel Computational Impact on Theory and Experiment
IT	information technology
LBNL	Lawrence Berkeley National Laboratory
LCF	Leadership Computing Facility
LLNL	Lawrence Livermore National Laboratory
LEED	Leadership in Energy and Environmental Design
LSMS	locally self-consistent multiple scattering
MD	molecular dynamics
MFDn	Many Fermion Dynamics, nuclear (nuclear shell model application)
MPI	message passing interface
MSM	multilevel summation method
MTTF	mean time to failure
MTTI	meant time to interrupt
NCAR	National Center for Atmospheric Research
NCCS	National Center for Computational Sciences
NCSU	North Carolina State University
NERSC	National Energy Research Scientific Computing Center
NICS	National Institute for Computational Sciences (University of Tennessee)
NOAA	National Oceanic and Atmospheric Administration
NRS	network request scheduler
OA	operational assessment
OMB	Office of Management and Budget
OLCF	Oak Ridge Leadership Computing Facility
ORISE	Oak Ridge Institute for Science and Education
ORNL	Oak Ridge National Laboratory
OS	operating system
OTF	Open Trace Format
PB	petabyte
PI	principal investigator
PNNL	Pacific Northwest National Laboratory
POP	Parallel Ocean Program
QoS	quality of service
RAMS	Research Alliance in Math and Science (ASCR)

RMP	risk management plan
RMTAP	Risk Management Techniques and Practice
RT	Request Tracker (ticket tracking software)
RUC	Resource Utilization Council (OLCF)
SA	system availability
SC	Office of Science (DOE)
SCaLeS	*A Science-Based Case for Large-Scale Simulation* (DOE SC report)
SciComp	Scientific Computing Group (OLCF)
SDN	Science Data Network
SNL	Sandia National Laboratories
SSM	storage system management (part of HPSS software)
SWC	Software Council (OLCF)
TB	terabyte
TechInt	Technology Integration Group (OLCF)
TUD	Technical University of Dresden
UAO	User Assistance and Outreach Group
VACET	Visualization and Analytics Center for Enabling Technologies
Vampir	Visualization and Analysis of MPI Resources (TUD)
WBS	work breakdown structure
WL	Wang-Landau (Monte Carlo method)
XGC	X-point included Guiding Center (code)—gyrokinetic particle-in-cell code designed for studying fusion plasmas
YTD	year to date

INTRODUCTION

Oak Ridge National Laboratory's (ORNL's) Cray XT5 supercomputer, Jaguar, kicked off the era of petascale scientific computing in 2008 with applications that sustained more than a thousand trillion floating point calculations per second—or 1 petaflop. Jaguar continues to grow even more powerful as it helps researchers broaden the boundaries of knowledge in virtually every domain of computational science, including weather and climate, nuclear energy, geosciences, combustion, bioenergy, fusion, and materials science. Their insights promise to broaden our knowledge in areas that are vitally important to the Department of Energy (DOE) and the nation as a whole, particularly energy assurance and climate change.

The science of the 21st century, however, will demand further revolutions in computing, supercomputers capable of a million trillion calculations a second—1 exaflop—and beyond. These systems will allow investigators to continue attacking global challenges through modeling and simulation and to unravel longstanding scientific questions. Creating such systems will also require new approaches to daunting challenges. High-performance systems of the future will need to be codesigned for scientific and engineering applications with best-in-class communications networks and data-management infrastructures and teams of skilled researchers able to take full advantage of these new resources.

The Oak Ridge Leadership Computing Facility (OLCF) provides the nation's most powerful open resource for capability computing, with a sustainable path that will maintain and extend national leadership for DOE's Office of Science (SC). The OLCF has engaged a world-class team to support petascale science and to take a dramatic step forward, fielding new capabilities for high-end science. This report highlights the successful delivery and operation of a petascale system and shows how the OLCF fosters application development teams, developing cutting-edge tools and resources for next-generation systems. Highlights from 2009–2010 include the following.

- A major system upgrade for Jaguar, from AMD Opteron "Barcelona" quad-core processors to AMD Opteron "Istanbul" six-core processors.
- Number 1 ranking on the Top500 list, the first time a DOE SC system has held the position of most powerful computer in the world, and overall top ranking in the HPC Challenge benchmarks (winning three of the four benchmarks), which reflects the balance of the system.
- Award-winning use: The 2009 Gordon Bell Prize winners used more than 223,000 of Jaguar's 224,000-plus available processing cores and achieved 1.84 petaflops using locally self-consistent multiple scattering (LSMS) to explore magnetic properties of materials with potential for radically enhancing a secure energy future. A Gordon Bell Prize finalist produced a sustained performance of 1.39 petaflops on 223,200 processors using NWChem, demonstrating proof-of-principle that the code can accurately model the weak bond in water.
- The production release of Spider, the world's largest and fastest Lustre file system, which is operating at 99.7% scheduled availability.
- Continued praise from users, with the high overall satisfaction ratings increasing annually from 3.7 in 2006 to 4.3 in 2009 (measured on a 1–5 scale in which 4 is "satisfied" and 5 is "very satisfied").

High-Performance Computing Facility Operational Assessment

Each year the DOE SC conducts an operational assessment (OA) of the performance of the OLCF and the other DOE high-performance computing facilities (HPCFs). This high-performance computing facility operational assessment (HPCFOA) is an SC programmatic management tool for evaluating the HPCFs' plans for providing high-performance computing and network resources as well as support to the scientific user base. Relevant information from the HPCFOA is used to respond to the Office of Management and Budget (OMB)–Office of the Chief Information Officer annual operational analysis data call for major information technology (IT) operations. A subset of the performance measures may also be included on the Information Technology Exhibit 300 for the facility. The HPCFOA focuses on the measurement of results and achievements in the following areas: customer support, business and strategic results, financial performance, and innovation. For each of the areas of interest, HPCFs respond to a series of "charge questions" using methodologies developed with the concurrence of and guidance from the federal program manager.

The charge questions and summary responses are as follows.

CHARGE QUESTION 1: Are the processes for supporting the customers, resolving problems, and communicating with key stakeholders effective?

OLCF RESPONSE: The OLCF continues to provide outstanding support for customers, exceeding the user satisfaction and problem resolution metrics.

- **Overall OLCF score on the user survey will be satisfactory (3.5/5.0) based on a statistically meaningful sample**. (OLCF 2009 survey rating is 4.3.)
- **Annual user survey results will show improvement in at least 1/2 of questions that scored below satisfactory (3.5) in previous period.** (No scores below 3.5 in the 2009 survey.)
- **80% of OLCF user problems will be addressed within three working days, by either resolving the problem or informing the user how the problem will be resolved.** (87% of problems were addressed within 3 working days.**)**

The overall user satisfaction rating has increased in each of the past 4 years, from 3.7 in 2006 to 4.1 in 2007, 4.2 in 2008, and 4.3 in 2009 (measured on a 1–5 scale in which 4 is "satisfied" and 5 is "very satisfied"). There were no results below 3.5 in either 2008 or 2009 (Section 1.1.1). OLCF staff members also ensure prompt response to users, exceeding the metric that 80% of user problems—"tickets"—be addressed within 3 working days. Between July 1, 2009, and June 30, 2010, 87% of tickets were addressed within 3 working days, with an average time to initial response of 31 minutes (Section 1.1.2).

The Scientific Computing Group (SciComp) liaisons and the center's other operational teams continue to shine as they work closely with users to do "whatever it takes" to ensure successful and effective use of the OLCF production resources, including the following.

- Assisting users to quickly and effectively identify code bugs and providing training in using leadership systems.
- Profiling code performance to isolate and resolve bottlenecks. In the case of GROMACS, a widely used molecular dynamics (MD) package (Section 1.2.1), improvements prompted by profiling the code are scheduled for public release, benefitting the MD community at large.

- Streamlining input/output (I/O) within codes such as RAPTOR, a massively parallel flow solver optimized for large eddy simulations, and the Parallel Ocean Program (POP), which enabled much larger and higher-resolution simulations than previously thought possible (Section 1.2.1).

The OLCF user support practice is to work directly with users to solve problems and communicate solutions to the broader community, to train users and students so they can participate in the evolution of codes to meet requirements for larger and larger simulations, and to proactively identify algorithm and run parameter improvements to most effectively make use of leadership systems. The OLCF communicates these practices and the subsequent accomplishments of staff and users through a wide variety of means, including a newly redesigned website, workshops and lecture series, conference calls, and media highlights.

CHARGE QUESTION 2: Is the OLCF maximizing resources consistent with its mission?

OLCF RESPONSE: The OLCF continues to deliver leadership computing resources to its users for cutting-edge simulations surpassing the following metrics.

- **Scheduled system availability: 85% in the first year of operation, 95% thereafter.**
- **Overall system availability: 80% in the first year of operation, 90% thereafter.**
- **Capability use: 35% of usage is from jobs using 20% or more of available cores.**

The center yielded its highest-ever overall user satisfaction rating even as it underwent a major computer system upgrade from AMD Opteron Barcelona quad-core processors to AMD Opteron Istanbul six-core processors. The primary production resources of the OLCF surpassed the targeted scheduled and overall availability metrics in calendar year (CY) 2009 to achieve the following.

- Cray XT5: 95.3% scheduled availability; 86.7% overall availability.
- Cray XT4: 97.4% scheduled availability; 94.0% overall availability.
- HPSS (High-Performance Storage System): 99.6% scheduled availability; 99.3% overall availability.
- Spider (External Parallel File Systems): 99.8% scheduled availability; 96.5% overall availability.

The metrics for CY 2010 year to date (YTD) meet or exceed the CY 2009 performance (see Section 2 for comprehensive metrics).

Through well-established committees and processes, the resources of the OLCF are effectively allocated through three programs:

- the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program (60% of available resources), now jointly operated by the Argonne Leadership Computing Facility (ALCF) and the OLCF and managed by the OLCF’s Julia C. White;
- the Advanced Scientific Computing Research Leadership Computing Challenge (ALCC) (up to 30%), managed by the DOE Advanced Scientific Computing Research (ASCR) Program Office; and
- the Director’s Discretionary (DD) program (10%), managed by James J. Hack, director of the National Center for Computational Sciences (NCCS).

Through the INCITE program, which identifies and grants large awards of time to leadership-scale science and engineering projects, awards increased from 470 million processor hours in CY 2009 to 950 million hours in CY 2010. The center continues to excel in supporting users and their application development as they graduate to larger and more diverse systems. INCITE usage in CY 2009 was 106%, and usage halfway through CY 2010 is on track at 57% through June 23, 2010. The effectiveness of the usage is just as important as the raw number of consumed compute hours.

The center tracks leadership usage as the total percentage of time used by jobs requesting 20% or more of the cores. At 36.4%, the OLCF meets the difficult target metric of 35% in CY 2010 YTD. We achieved this even though the physical number of cores increased by 50% as part of the six-core Opteron upgrade.

CHARGE QUESTION 3: Is the OLCF meeting Department of Energy Strategic Goals 3.1 and 3.2?

OLCF RESPONSE: As the first center to run multiple production application codes at a sustained petaflop, the OLCF continues to provide researchers with an effective leadership system that enables simulations at an unprecedented scale, both in terms of the physical parameters that can be incorporated and resolutions that can be reached. Detailed quantitative answers to fundamental questions regarding microscopic interactions in water and in magnetic materials were produced by two of the most computationally intensive simulations ever performed. In each of these cases, the power of the world's first petascale computer for open science was not only useful, but essential (Section 3 sidebars).

In addition, the OLCF continues to provide core resources for community climate research, from simulations of decades- to centuries-long climate events to code development that allow applications to provide results with an accuracy sufficient for policy discussions and decision making. Nuclear reactor studies that would be untenable in an experimental setting are now possible with Jaguar. Increasingly detailed simulations of nuclear reactor cores are using 70% or more of the Cray XT to produce more realistic models and, ultimately, reactor designs. These leadership research investigations and others can be found in Section 3. In 2009 the OLCF initiated an Industrial Partnerships Program to promote usage of high performance computing across U.S. industry. For more details about this program and recent success stories, see Section 5.9.

CHARGE QUESTION 4: How well is the program executing to the cost baseline pre-established during the previous year's Budget Deep Dive? Explain major discrepancies.

OLCF RESPONSE: The total cost for fiscal year (FY) 2010 was $96,114K. Of this, 21.3% was spent on effort, 43.7% on lease payments, 10% on center charges (utilities), 10.1% on computer system maintenance, and 14.9% on other costs. The OLCF carefully managed costs in FY 2010 to execute the FY 2010 OLCF operational requirements (Section 4) and meet the targeted system availability and number of hours delivered. The financial status of the OLCF is monitored daily by the OLCF finance officer and no less than monthly by OLCF management.

CHARGE QUESTION 5: What innovations have been implemented that have improved the OLCF's operations?

OLCF RESPONSE: The OLCF has developed and implemented a series of high-impact innovations this past year to support the leading edge scientific breakthroughs being pioneered on OLCF platforms and to advance high-performance computing (HPC).

<u>Center-Wide File System</u>. In FY 2010 the OLCF released "Spider," the world's largest Lustre file system (Section 5.1). This center-wide parallel file system demonstrated more than 240 GB per second (GB/s) of aggregate bandwidth on the full storage system, with several performance improvements allowing the file system to achieve more than 91% of the raw disk bandwidth (267 GB/s). Performance on the DataDirect Networks (DDN) S2A9900 hardware was nearly doubled by eliminating concurrent 5 ms pauses from

each storage request. Further innovations have also reduced single points of failure, maximizing availability (99.7% scheduled availability).

ORNL's development efforts on the Spider file system have had a significant impact on Cray's I/O plans and have greatly influenced the scope of I/O solutions Cray now offers its customers, ultimately resulting in a new product line available to Cray customers and the development of next-generation file system technologies.

<u>Application Tracing, Analysis, and Optimization</u>. The OLCF recognizes that leadership computing requires the development of tools to support users on the world's largest compute system. Applications that are run at this unprecedented scale will produce errors not previously encountered. The OLCF provides an integrated system monitoring framework, with customized tools, that collects system events and aggregates performance data, allowing for detailed analyses of system failures and aberrant performance and root cause analysis of system interrupts. Innovations developed by the OLCF include the following.

- MDSTrace, which allows staff to identify applications that create substantial load on our parallel I/O environment in an automated fashion. We anticipate releasing this tool to the broader community as an open source software project (Section 5.2.1).
- A collaboration with North Carolina State University (NCSU) to customize ScalaTrace for the Cray XT environment. These extensions allow any site using Cray XT systems to leverage the tool to identify the root cause of application performance degradation (Section 5.2.1).
- DDNTool, a flexible monitoring tool for our storage system infrastructure (Section 5.2.2).

<u>Improving Application Performance Through Operating System Scalability</u>. In 2009, OLCF systems analysts observed a high degree of variability in application performance on the Cray XT5 system. During the discovery process, the OLCF staff found higher levels of operating system (OS) activity on the Cray XT5 than on other HPC platforms, impacting performance on some applications. In response, the OLCF initiated a collaborative effort with Cray to address the problem at leadership scale. The Cray-OLCF team weighed various options to address the problem and determined that a reduced noise kernel would be the most practical, providing a common Linux environment coupled with the scalability requirements of leadership computing. See Section 5.7 for details. Innovations resulting from this collaboration include the following.

- The ability to isolate OS services to a single processing core, leaving other processing cores dedicated to application tasks. This feature is selectable by the application at job submission.
- Performance improvements of up to two orders of magnitude in large-scale collective communication.
- Improved performance for communication-bound applications such as POP (30% performance improvement and improved strong scaling).

A rich variety of activities have resulted in tools and enhanced support in other areas, including parallel data tools; integrated end-to-end scientific data management, including the widely used adaptable I/O system (ADIOS); centralized maintenance of software; HPSS; and scalable debugging and performance tools, to name a few. These are described in greater detail in Section 5 of the report.

CHARGE QUESTION 6: Is the OLCF 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. The risk management plan includes:

- identifying and analyzing potential risks,
- ensuring that risk issues are discovered and understood early on,
- ensuring that mitigation plans are prepared and implemented, and
- developing budgets with consideration of risk.

The OLCF currently tracks 25 operational risks. Three are categorized as "High," five "Medium," and the remainder a "Low" level of risk. At periodic risk reviews, weekly staff meetings, and ad hoc discussions, the OLCF management team focuses attention on the high and the moderate risks while keeping an eye on the low risks, which may change over time.

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).

CHARGE QUESTION 7: Does the OLCF have a valid cyber security plan and authority to operate?

OLCF RESPONSE: The OLCF maintains an effective and up-to-date cyber security program plan (CSPP). DOE awarded the facility continued Authority To Operate on September 22, 2009, after an updated CSPP with enhancements in the user identity proofing process was submitted. The current accreditation of the ORNL cyber enclave expires June 21, 2011.

RESPONSES TO 2009 OPERATIONAL ASSESSMENT COMMENTS AND RECOMMENDATIONS

Comment/Recommendation: The centralized file system availability metrics should be presented in the future.

OLCF Response: Based on the recommendations of the 2009 committee, new metrics have been introduced. The scheduled and overall availability of the centralized file system should be at least as high as that of the production system; therefore, the metrics we propose for tracking the centralized file system are 95% scheduled availability and 90% overall availability (Section 2).

Comment/Recommendation: Risk mitigation strategies and costs should be documented. Continue to assess/monitor staff planning to adapt to changing user needs and/or funding impacts and link this to your risk management.

OLCF Response: The OLCF risk register has all risks and mitigation strategies documented for High and Medium risks. The costs for this are included in the operational budget as part of the cost of doing each task.

Staff planning is based on the needs of our user community as well as the requirements to effectively run the computer systems. The staffing budget is linked to the overall budget. The risk register includes funding risks that impact our staffing levels.

Comment/Recommendation: Performance metrics to determine the true impact of various capability calculations should be continually improved and refined.

OLCF Response: The OLCF currently follows the recommendation in the 2007 report[*] of the Advanced Scientific Computing Advisory Committee (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). It is exceedingly difficult to formulate metrics for assessing the impact of capability simulations. As noted by the ASCAC Petascale Metrics Panel, breakthroughs are "an immeasurable goal" with no obvious metrics that would predict "discoveries that occur on the leading edge of fundamental science."[†]

This issue (i.e., how to formulate meaningful metrics for impact) is also the subject of considerable interest and debate among information science practitioners, and a consistent view has yet to emerge regarding the efficacy, or even the desirability, of such metrics.[‡] Important results require time to be understood and appreciated by other workers, but it is this eventual integration into subsequent work and thought that represents the full measure of scientific impact. To that end and in keeping with the 2008 ASCAC Committee of Visitors (COV) recommendation, "in approximately five years [e.g.,

---

[*] Panel recommendations can be found in the full report of the committee, *Advanced Scientific Computing Advisory Committee Petascale Metrics Report*, 28 February 2007, available at http://www.er.doe.gov/ascr/ascac/Reports/PetascaleMetricsReport.pdf.

[†] Ibid., pp. 5 and 19.

[‡] Numerous articles have been written on the topic; the following give a good overview of the issues: K. L. Reed, "Citation analysis of faculty publication: beyond Science Citation Index and Social Science Citation Index," *Bull. Med. Libr. Assoc.*, 83(4), pp. 503–508 (1995); Lowell L. Hargens and Howard Schuman, "Citation counts and social comparisons: scientists' use and evaluation of citation index data," *Social Sci. Res.*, 19(3), pp. 205–221 (1990); A. Sidiropoulos, D. Katsaros, and Y. Manolopoulos, "Generalized Hirsch h-index for disclosing latent facts in citation works," *Scientometrics*, 72(2), pp. 258–280 (2007).

circa 2012–2013] a formal review panel [will] be convened to assess the impact of the INCITE program."[*] As joint administrator of the INCITE program, the OLCF will ensure that this review is informed by the full measure of scientific discovery and engineering progress realized through the capability-class simulations of our users.

Comment/Recommendation: The down-select process of the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program should also be continually improved and refined.

OLCF Response: Since 2009 we have implemented several of the remaining recommendations for the INCITE program made by the 2008 ASCAC COV. The total number of INCITE awards increased incrementally from 66 in 2009 to 69 in 2010, and we anticipate the total number of awarded projects to remain flat in 2011, in keeping with the committee's recommendation that the number of projects be kept low. Based on input from panels composed of leaders in the field, only the most highly rated proposals received awards of time last year: The acceptance rate for new INCITE proposals was about 40% for the 2010 award period. As advised by the COV, renewals are held to a standard below which projects are rejected. In last year's call for proposals, 84% of the renewal requests received awards of time to continue leadership-level computing (Section 3.4).

[*]Advanced Scientific Computing advisory Committee, Committee of Visitors, *INCITE Report,* 2008, available at http://www.er.doe.gov/ascr/ascac/Reports/INCITECOVReportAug08.pdf.

FY 2010 METRICS

1. CUSTOMER SUPPORT

CHARGE QUESTION 1: Are the processes for supporting the customers, resolving problems, and communicating with key stakeholders effective?

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; comprehensive training programs and user workshops; and tools to reach and train the next generation of computer scientists (Figure 1.1).

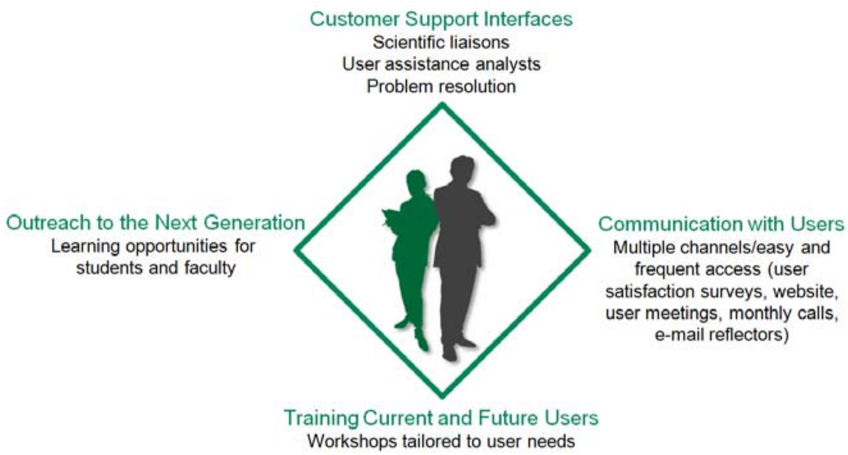

Figure 1.1. Customer support and outreach model.

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 Oak Ridge facility is expanding this outreach through an internship program for science writers: by working alongside senior science writers at the facility and with computational researchers, these interns gain a more thorough understanding of the impact of leadership computing, and this is translated into more insightful news stories as these students transition to other media outlets (Section 1.4).

The role that communication plays in everything we do cannot be emphasized too strongly. The center comprises many different, complementary groups, but the groups are physically located in close proximity in the OLCF and communicate well to solve problems in a collegial manner (see examples in Sections 1.1, 1.2, and 5).

Annual user surveys, with metrics agreed upon by the OLCF management, OA review panel, and DOE program managers, provide information on how we are doing and input for future planning. Survey results have shown consistently high satisfaction rates, and this year was no exception, with 90% of responders indicating they were satisfied or very satisfied with OLCF services. There has been a year-to-year increase in overall satisfaction with the OLCF from 3.7 in 2006 to 4.1 in 2007, 4.2 in 2008, and 4.3 in 2009, out of a possible score of 5.

Day-to-day user problems and query resolution are managed with the popular Request Tracker software (RT). A user assistance analyst from the User Assistance and Outreach Group (UAO) is assigned to each query until it is resolved. Our goal, which we have consistently met, has been to **address** at least **80% of tickets within 3 days** and to at least touch base with users within 2 hours. When asked about the speed of initial contact and quality of response, a majority of users were satisfied or very satisfied (89% and 88%, respectively). Several users specifically commented in the survey on the OLCF response time, using words such as “fast,” “immediate,” and “quick.” The following are characteristic responses.

- “Immediate response from the user help.”
- “Knowledgeable staff which are very quick and helpful in solving problems.”

Scientific liaisons are a unique OLCF response to high-performance scientific computing problems faced by users. The OLCF recognized early on that users of HPCFs 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 responded with two complementary OLCF user support vehicles: UAO and SciComp, which includes the scientific and visualization liaisons.

Multiple channels have been developed for communicating with users, from the formal (user surveys) to the informal (day-to-day interactions with users). Liaisons and UAO members solicit input from users during formal and informal interactions and act as advocates, both informally and through formal mechanisms such as the Resource Utilization Council (RUC). As one example of our response to user feedback, we have renovated our website to make it more accessible and responsive to our users.

The OLCF offers many training and educational opportunities throughout the year for both current facility users and the next generation of HPC users (Sections 1.4 and 3.3).

1.1 EFFECTIVE CUSTOMER SUPPORT

A multifaceted approach is used to measure the effectiveness of the OLCF customer support model: a yearly survey measures customer satisfaction in key areas, a query management system ensures all queries are responded to in a timely manner, and OLCF staff members solicit feedback directly from stakeholders through various formal and informal interactions.

1.1.1 User Survey

The OLCF conducts an annual survey of all users to solicit feedback on the quality of our customer service and computational resources. The survey is conducted by an independent third party, the Oak Ridge Institute for Science and Education (ORISE), using questions developed by the OLCF in collaboration with the DOE OLCF program manager and with input provided by ORISE. The surveys, which contain 50 questions, are made available online to all individuals with active accounts (698 this year, excluding OLCF staff and vendors); periodic reminders are sent to nonresponders. Survey results are validated using a streamlined version of the Delphi Technique, a set of guidelines for remote gathering of information from experts.

This year, 261 out of a total of 698 users responded to the survey for a response rate of 37%. While the rate is lower than for last year, the total number of responders actually increased from 226 to 261 due to a higher number of users (Table 1.1).

Table 1.1. Overall Characteristics of Respondents

	2008 Survey	2009 Survey
Total Responders	226 (48%)	261 (37%)
New Users (OLCF user <1 year)	41%	29%
OLCF User for 1–2 years	27%	36%
OLCF User >2 years	32%	35%
Used User Assistance Center at least 1 time	82%	74%

Users are 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 our program manager define **3.5** to be **satisfactory**. The actual scores in Table 1.2 and comments in Table 1.3 indicate that users are very satisfied with OLCF customer service and computational resources.

Table 1.2. Average Satisfaction Rates for Key Indicators

Indicator	2008 Survey	2009 Survey	Difference (%)
Overall satisfaction with the OLCF	4.2	4.3	+2.4
Resolution of queries by the UAO	4.2	4.4	+4.8
Overall system performance of the XT5	NA	4.1	—

Table 1.3. User Comments from the 2009 Survey

"Of the three super computer sites I work at, you have the most responsive consulting staff by far."
"The OLCF provides first rate computing resources. In addition, I have found the center to be extremely welcoming during my visits to ORNL."
"I like the quick and clear answers given via e-mail. And of course I like the fast supercomputers."
"The machines are very powerful and well maintained. The supporting staff is responsive."
"Systems and supporting software/hardware resources are complete and bug free. All systems are intuitively easy to use and port codes to."
"The OLCF staff do an exceptional job of providing timely responses to the questions on everything from login problems to software issues. They have been a tremendous asset in helping me to port our in-house docking code to the Jaguar platform."
"Obviously the systems available are the best in the world. But the user support has also been wonderful. Anytime I have had an issue it is resolved very quickly."
"The XT5 and XT4 systems are down too often."
"You should have introductory workshops for beginners."

Each year the OLCF works to show improvement in at least half of the questions that scored below satisfactory (3.5) in the previous year's survey. **All questions scored above 3.5 in both 2008 and 2009**,[*] but the OLCF continues to explore and introduce new products (for example the redesigned website) and more effective and efficient processes to promote user satisfaction.

Because the surveys are one of the tools we use to continually improve operations, users are also asked a few open-ended questions to solicit feedback on our strengths and specific areas for improvement. In response to an open-ended question about the best qualities of the OLCF, thematic analysis of user responses identified the following as the top three:

- great staff and support (40% of responses),
- powerful/fast machines (34% of responses), and
- computational resources [29% of responses (overlap with "powerful/fast machines"[†])].

In the 2009 survey, the following areas for improvement were cited the most frequently.

- Reliability (21%)
- Stability (17%)
- File System (13%)
- User support (9%)

The reliability and stability of OLCF systems increased this year (Section 2). While the 9% of responders to last year's survey who suggested user support as an area for improvement is small (roughly 12 responders), the OLCF takes all comments and suggestions seriously and has implemented a number of changes to address these comments and those received through other, less formal channels.

- One user specifically requested better communication of changes to modules and compilers before the changes occur. UAO worked with the Software Council (SWC) discussed in Section 2 to address this issue. The SWC agreed to reduce the frequency with which changes were made and to perform all module and compiler changes on a quarterly basis whenever possible. In addition, all changes of this nature are announced in the weekly e-mail before they are made.
- A few survey respondents requested improvements to the center's website. The website has been completely redesigned (Section 1.3).
- Some users requested the ability to keep data in scratch space longer. UAO addressed this issue by creating project areas that are not purged for all users. Users are now able to keep data in the scratch area until they are ready to transfer the data elsewhere.
- Users noted a need for improvements to the debugging process. Much work has been done in this area over the past year, and major accomplishments and innovations are discussed in Section 5.
- One user requested introductory training workshops (Table 1.3). During the last spring workshop (May 2010), OLCF addressed the need for different levels of training by providing both a beginner and an advanced track. This allowed us to accommodate both new and returning users.

---

[*] The lowest rating received in the 2009 survey, a 3.6, was in response to question 34: "Overall rate your satisfaction with the following aspects for the Development Cluster (Smoky) platform." Respondents rated the scratch disk size/performance "aspect" on Smoky 3.6.

[†] In looking at the responses, users differentiated between hardware (powerful/fast machines) and other computational resources such as software and storage.

- A few users requested more online training materials. In response, UAO developed additional online training documents, including the following.
  - Introduction to Parallel Computing with MPI
  - Using Parallel I/O
  - PGI and Cray Compiler Optimization
  - MPI Optimization and Tips
  - Introduction to the CrayPAT

1.1.2 Problem Resolution

The OLCF uses RT to track queries and ensure that response goals are not missed. In addition, the software collates statistics on tickets issued, turnaround times, etc., to produce weekly reports so that we can track patterns to address anomalous behaviors before they have an impact on additional users. The OLCF issued more than 2,700 tickets in response to user queries between July 1, 2009, and June 30, 2010 (Figure 1.2). The team met both the response time and the resolution time metrics:

- the average response time for a query was 31 minutes and
- 87% of queries were addressed within 3 working days.

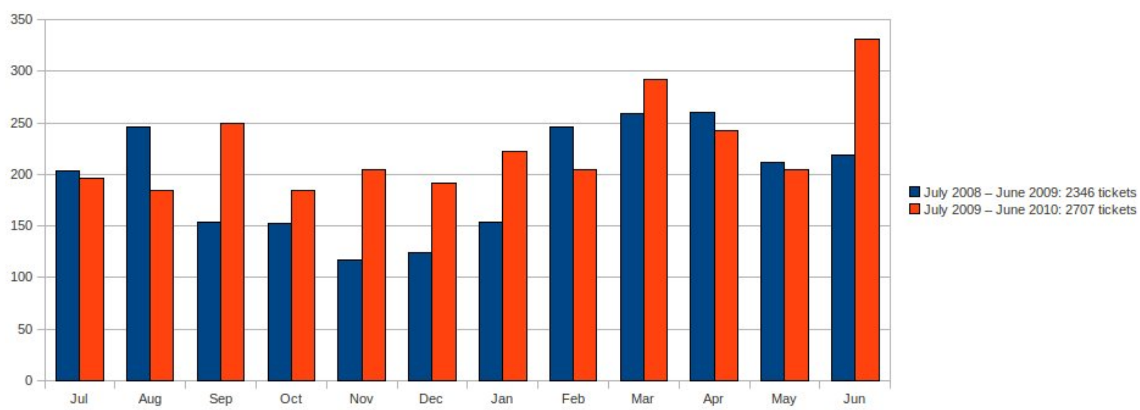

Figure 1.2. Number of tickets issued per month.

Tailoring services to the needs of our users is a key element in the success of the OLCF. 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. 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.

UAO holds routine ticket report meetings during which members share information on issues, solutions, etc. As a result, we have further streamlined the issue resolution process by providing both lessons learned on simple problems and development opportunities for newer staff members.

One change resulting from "ticket meetings" involves the process for activating the SecureID password token, which users need for logging into the system. Users previously needed to call the helpline to activate the token. After discussion in a UAO meeting, however, it was determined that this could be accomplished via e-mail, eliminating an extra step and the accompanying wait.

UAO members routinely provide the following types of support to our users.

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

Installing third-party applications

Because the market is relatively small and there is very little commercially available software for supercomputers, much of the software that runs on our systems is open source. Thus when a new version comes out, it falls to the OLCF to obtain, build, and test it. From July 2009 through June 2010, the OLCF installed 52 applications and made 47 default version changes.

One such package is the Earth System Modeling Framework (ESMF) developed at the National Center for Atmospheric Research (NCAR). This package is used for climate simulations, weather prediction, and similar computationally demanding applications that run well only on supercomputers such as those at the OLCF. INCITE projects based at the National Oceanic and Atmospheric Administration (NOAA) and NCAR use ESMF on OLCF systems. At one time each of our ESMF users had to download a version of the software and build it independently—leading to duplication of effort and, in some cases, errors that bogged down the system. In response, SciComp member Ilene Carpenter worked with the ESMF developers (now based at NOAA) to build and install the latest version on our systems, thereby freeing individual users from having to build and (potentially) debug this complex software package. As a result, meteorologists, earth scientists, and others working with complex climatological data on our systems can focus on their own applications with confidence while using a version of the ESMF library that has been tested by the ESMF developers.

1.2 USER SUPPORT—WHATEVER IT TAKES

1.2.1 Scientific Liaisons

In preparing for the OLCF, planners at ORNL envisioned the perfect environment for helping HPC users address the daunting challenges posed by a unique supercomputer such as Jaguar. These problems, they quickly understood, would completely overwhelm a traditional help desk and would even be too much for UAO. In response, they created the OLCF's much-duplicated SciComp liaison program: experts in user support—including PhD-level liaisons from fields such as astrophysics, climatology, applied mathematics, and combustion engineering—who are also specialists in developing code and optimizing it for the OLCF systems. The liaisons are, in a sense, intermediaries between scientists and machines. They are embedded on project teams and provide needed support, from debugging code, to advocating for users, to full science and project collaboration.

Benefits from the liaison program are broad and far reaching. Information gained by one liaison is shared with other teams. In addition, liaisons act as advocates for the principal investigators (PIs) and projects they support, communicating, for instance, with computer, software, and tools vendors and with other OLCF groups. Working as a team, the liaisons leverage one another's expertise to ensure projects run in an efficient, effective way. For example, particle-tracking techniques used in combustion simulations are applicable to fusion or astrophysics efforts. Liaisons also share information on the peculiarities of the machines—"things to watch out for," in the words of one of the liaisons. It's a win-win: Projects get consultants with scientific expertise who are also experts at running simulations on the world's biggest supercomputer, and the OLCF gets more efficient codes that run more effectively.

This approach also provides a nurturing, exhilarating environment for scientists. Markus Eisenbach, OLCF liaison and 2009 Gordon Bell Prize winner, worked in the Materials Science and Technology Division before joining the OLCF 3 years ago, but he was always involved on "the computing side," as he says. What attracted him to the OLCF was the broad spectrum of opportunities: While supporting users, he could still do his own original research and maintain contacts with the larger scientific community.

The fact that Jaguar has been said to be "easy to use while delivering terrific results" can be attributed in no small part to the liaison program.*

Today, OLCF liaison support encompasses a range of activities—"whatever it takes," as one manager is fond of saying—including the following.

- Improving performance and scalability of project application software.
- Assisting in redesign, development, and implementation of strategies that increase effective use of OLCF resources.
- Implementing scalable algorithm choices and library-based solutions.
- Providing an advocacy interface to OLCF resource decisions, including the RUC.
- Performance modeling such as anticipating the impact of upgrades and fine-tuning applications for maximum efficiency.
- Scaling applications to make effective use of the OLCF's petascale resources.
- Assisting with code development and algorithms.

Improving performance and scalability

A "misbehaving" nuclear physics code this past year provided an excellent example of OLCF intergroup communications to problem solve and provide leadership computing support. In this particular case, periods of extreme slowdown on the Cray XT system were adversely affecting users and their application runs—something usually indicative of improper use of the system by one or more applications. Three groups, TechInt, HPC Operations, and SciComp, worked collaboratively to identify the issue and correct the problem. HPC Operations and TechInt used monitoring tools developed by TechInt (Section 5) and found the system substantially degraded in performance every time certain codes ran. Using OLCF tracing software developed by TechInt (Section 5), the groups identified specific user application runs that corresponded with performance slowdowns.

One of the applications involved was vital for the "Computational Nuclear Structure" INCITE project for which Hai Ah Nam is the liaison. Because of her background in physics and knowledge of Fortran, she was readily able to compare the trace information to the user's code, identify the specific problem area (the I/O), and modify the code to run more efficiently. Nam's modification to the code's I/O not only enabled the code to run more efficiently without degrading system performance, it cut the code's run time in half and corrected a race condition that had impacted the correctness of results [because all of the processors (48 K) were doing an open-seek on one file to populate 13 variables, the values of the variables were not set on all processors, and so the calculation gave incorrect results]. Nicolas Schunck, the project principal investigator (PI), was impressed with the speed with which the problem was detected and resolved. Further, he says that "The patch that was given to me made the entire code far more

*National Center for Computational Science, *Leadership Computing for Science: DOE Delivers Petascale Systems in 2008*, NCCS Annual Report 2008, ORNL 2009-G00827/JCP (Oak Ridge, Tennessee, Oak Ridge National Laboratory, 2009); available at http://www.nccs.gov/media-center/nccs-reports.

scalable and actually reduced the time of execution significantly. This solution has also been very beneficial for other projects I am involved with, which are also using similar I/O models."

Nam also taught a nuclear physics graduate student working on this project to use profiling and debugging tools [Visualization and Analysis of MPI Resources (Vampir) and Allinea's Distributed Debugging Tool (DDT)] to reduce his time-to-solution in new code development. Using DDT to step through the code, he was able to pinpoint the exact location of errors in a single trial rather than the multiple iterations that would have been required to find the error using only print statements (which far too many scientists depend on for debugging). (Note: Vampir and DDT training are an outgrowth of other projects discussed in more detail in Section 5.)

As this and the following examples demonstrate, I/O is still a major bottleneck in many codes.

Ocean I/O. Frank Bryan of NCAR discovered the challenge of I/O when he tried to run POP on the OLCF systems. POP, a popular stand-alone ocean model used in climate modeling (Figure 1.3), is the ocean component of the Community Climate System Model (CCSM), one of the most widely used general circulation models for climate research. Like many climate codes, POP does not scale well to large processor counts. In the past, POP was therefore run at low resolution on a small number of processors; however, Frank needed to run POP at high resolution (using more processors) to get more detail. SciComp's Jim Rosinski determined that I/O was the main problem slowing down the code; all I/O was going through one processor. He eliminated the problem by implementing a parallel I/O layer directing output to 42 processors instead of one, and in the process also accelerated I/O by a factor of 20.

Figure 1.3. Visualization of ocean temperature at 5 meters from POP simulation.

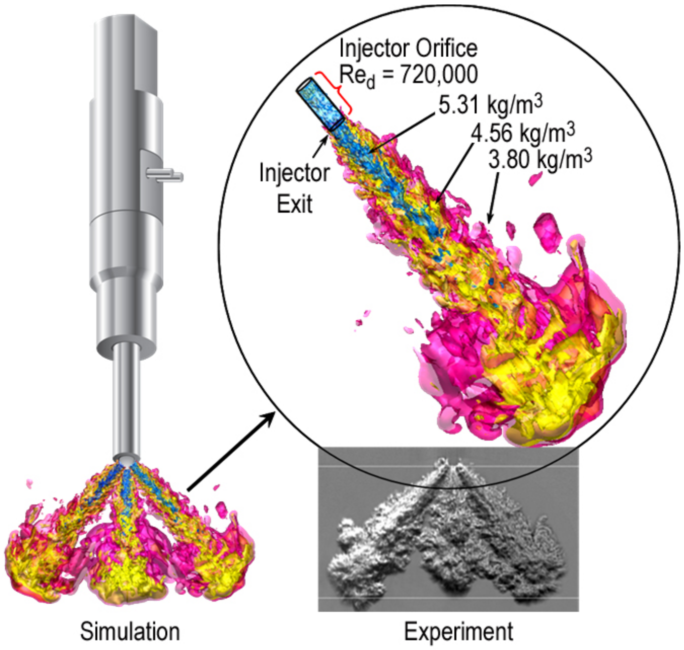

Figure 1.4. Large eddy simulation of high pressure injection processes associated with internal combustion engines validated against experiments.

Combustion Modeling. Modeling complex combustion processes is key to developing the next generation of internal combustion engines and other propulsion and power devices, including those using both conventional and advanced alternative fuels (Figure 1.4). Combustion processes, however, are highly complex; as a result, combustion simulations are inherently huge, requiring a high-performance computer like Jaguar. RAPTOR, a massively parallel flow solver optimized for treating turbulent combustion processes using large eddy simulation techniques, is currently one of the leading candidates for detailed simulations in engines. Researchers from Sandia National Laboratories' (SNL's) Combustion Research Facility have been running RAPTOR and a companion numerical simulation technique on Jaguar through a multiyear INCITE project allocation. Even though RAPTOR had exhibited strong scaling attributes in the past, it encountered I/O problems when it was ported to 47,000 processors on Jaguar; the code's I/O was based on smaller systems and did not work at petascale. The program was taking too long to write data, which got the PI's attention, and it was affecting system performance, which brought it to the attention of TechInt and the combustion project liaison, Ramanan Sankaran. For a short-term, immediate

Scaling up to petascale—and beyond

To choose wisely in picking future HPC systems, one must understand how key applications are likely to perform on such hardware. As a result, SciComp has begun characterizing the performance of major science applications in depth. The performance modeling effort attempts to predict how science applications will perform on future platforms by quantifying each application's utilization of different existing hardware subsystems (e.g., CPU, memory, communication network, and I/O). The application is executed under different user settings, for example differing problem sizes and core counts. The resulting performance data are used to infer the relative times spent in these subsystems and, thereby, to project how the code will perform on future computer hardware. The process may then improve application performance on the subsystems of the new hardware. This analysis also reveals the hardware "pressure points" for each application, that is, which hardware subsystems require most of the application runtime. By extension, it will also inform the priority we should give to the priority for each of these hardware subsystems in future HPC procurements.

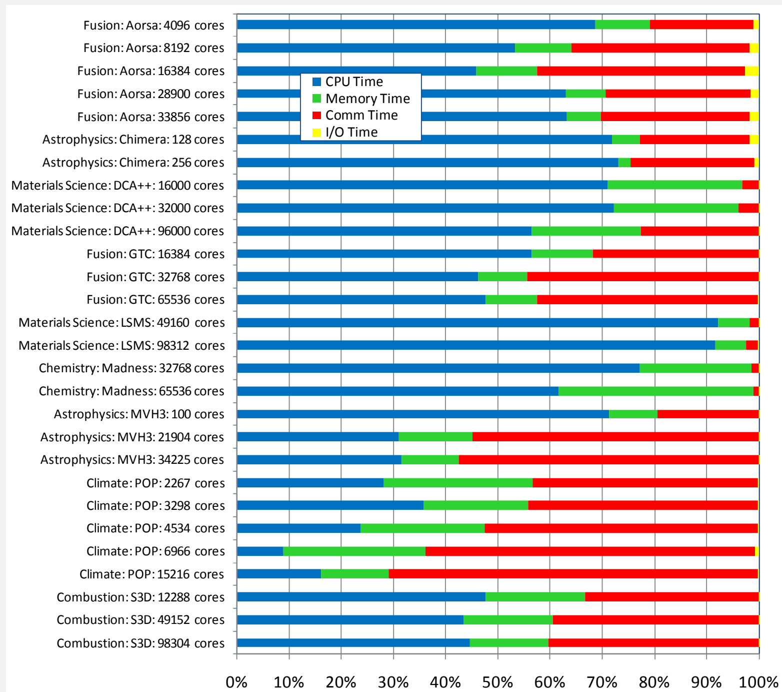

Application runtime by hardware subsystem.

Here are a few of the insights this process yielded that we can use in future planning.

1. Though memory, communication, and I/O are challenging and require substantial compute time, much runtime for many applications is still spent in CPU-bound work, implying that more powerful processors such as GPU accelerators could provide significant performance gains.
2. The communication interconnect is still very important for some science areas (e.g., climate), and its performance must not be neglected for future systems.
3. Some applications spend significant time in memory operations, suggesting the need for better memory-centric algorithms to hide memory latency.

This methodology for predicting performance has already played a part in the AMD Opteron Istanbul processor upgrade for the Jaguar platform, and results of these studies have been published in the latest OLCF requirements document, *Preparing for Exascale: ORNL Leadership Computing Facility Application Requirements and Strategy* (http://www.nccs.gov/wp-content/media/nccs_reports/olcf-requirements.pdf).

solution, Sankaran worked with the PI to reduce the number of files being written by aggregating and separating the data into multiple directories instead of one, easing the I/O impact. This allowed the team to proceed with planned simulations, avoiding the I/O hurdle and alleviating the impact on the system. With the immediate system stress resolved, Ramanan was free to work on a more elegant solution, which will be completed this year. He will be using the “multiblock” nature of the code to aggregate and write the data to disk. This will yield two benefits: (1) the original 47,000 blocks or files that were overloading the system will be aggregated and (2) the bigger data chunks will be easier for the PI to manage and analyze.

Bioenergy Bottlenecks. Overcoming the resistance of plant cell walls to hydrolysis is a major technological challenge for developing cellulosic bioethanol, a U.S. energy policy goal. To solve this problem is to allow vehicles—figuratively speaking—to run on straw . . . and to promote greater energy security. In a project to better understand this resistance, Jeremy Smith uses Jaguar to run highly parallelized computer simulations with the GROMACS MD code. MD simulations of biomolecular systems can easily run for several months—even on Cray XT systems like Jaguar.

In an effort to speed up the process, Arnold Tharrington, a SciComp biophysics expert, has been modeling GROMACS performance to identify bottlenecks in the code. One of the biggest, in terms of scalability in MD codes such as GROMACS, is the computation of long-range electrostatic interactions. Tharrington was unaware of the severity of the problem—especially at scale—until he and student interns began profiling code performance. As a result, the OLCF is spearheading a multiorganizational team that includes representatives from universities, industry (Cray), and other national laboratories to develop a scalable alternative known as the “multilevel summation method” (MSM). Tharrington, the task lead for developing the MSM library, said it was not difficult to get other people to collaborate, because the problem is large and the codes involved are important and pervasive. The scalable method the team is working on will be general enough for use in other MD codes as well, allowing them to overcome the bottleneck that plagues all MD codes. The method is scheduled to be put into a library that other groups can access, making a major contribution not only to the biosciences but to the entire HPC community.

Assisting with code development and algorithms

Nuclear technology is a necessary part of any current U.S. energy security strategy, so the Denovo code for radiation transport being developed by ORNL’s Tom Evans and colleagues (Figure 1.5) has great relevance. While all such codes are highly complex, more than 90% of the time spent in typical Denovo runs is devoted to the three-dimensional (3-D) sweep algorithm, a key component of the code that calculates radiation flow in the reactor. Evans and OLCF liaison Wayne Joubert were familiar with the algorithm and knew that it would be a bottleneck in any code using it. Therefore, when Evans was allocated time on Jaguar through the INCITE program to develop highly accurate spatial resolutions to validate the code, they knew it was time to look at optimizing it as well. Joubert is working to

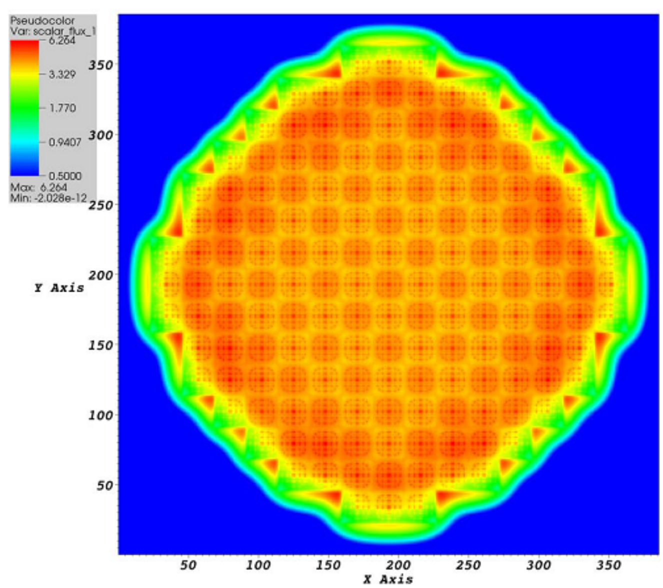

Figure 1.5. Thermal flux from Denovo reactor core simulation.

improve the sweep algorithm by developing and implementing innovative parallel optimizations for GPU processors. The GPU version of the 3-D sweep code is expected to perform faster than the current CPU-only version, substantially cutting the run time of Denovo. Because of the algorithm's computational expense, it was a good candidate for a GPU machine; however, the algorithmic changes made to enable GPU performance should result in better performance on conventional multicore processors as well. In addition, since the 3-D sweep algorithm is ubiquitous in radiation transport codes, the changes Evans and Joubert are working on will potentially have much broader applications. "Working with Wayne Joubert in NCCS, ORNL staff have advanced the state of Denovo's solvers in a significant way," Evans said. "NCCS has developed high-performance sweep kernels for next-generation GPU-accelerated computing hardware. Concurrently, researchers in the ORNL Nuclear Science and Technology Division have developed new parallel solution algorithms that can take advantage of this low level kernel. This parsing of work to staff with the correct domain expertise has resulted in a fruitful collaboration that has produced a next-generation computational transport code."

1.2.2 Visualization Liaisons

Most projects are assigned a visualization liaison in addition to a primary scientific liaison in order to maximize opportunities for success on the leadership computing resources. This approach stems from the recognition that scientific discovery relies on more than just volume of data. The ultimate goal is to make sense of the data. In fact, OLCF visualization scientists do more than strengthen a project's data analysis and help illuminate project results; in many cases they also help in detecting and fixing problems. In addition to customary visualization support services, OLCF visualization experts frequently find themselves developing custom software and algorithms to address the unique challenges of the users.

Writing custom visualization tools and algorithms

The visualization of vector fields is one of the more complex areas of scientific visualization. For example, the analysis of the fluid flow that governs natural phenomena on scales from the smallest (e.g., Rayleigh-Taylor mixing of fluids) to the largest (e.g., supernovae explosions) relies crucially on visualization to elucidate the patterns exhibited by flows and the dynamical aspects driving them. One common technique for visualizing these effects is through the use of streamlines, where a set of massless particles are advected through the vector field. Until recently, however, no scalable parallel methods existed for vector field analysis on petascale data sets. As part of DOE's Scientific Discovery through Advanced Computing Visualization and Analytics Center for Enabling Technologies (VACET) program, the OLCF's Sean Ahern and Dave Pugmire led a team including researchers from LBNL and the University of California–Davis that developed an algorithm for parallel computation of streamlines to help scientists analyze the vector fields being computed. As with many VACET projects, the effort was groundbreaking in that no parallel methods existed previously. Aiming for broad applicability, the team tested several prototype algorithms and settled on a hybrid method that achieves good performance over differing use cases. The tool has already been used at ORNL for reactor (Figure 1.6), astrophysics, and fusion simulations. In the case of the astrophysics simulations, the team was, for the first time, able to see streamlines through the magnetic field in a core-collapse supernova. The streamline tool has been deployed to VisIt for broad user community availability. Using this foundation, the team is now expanding its work to include

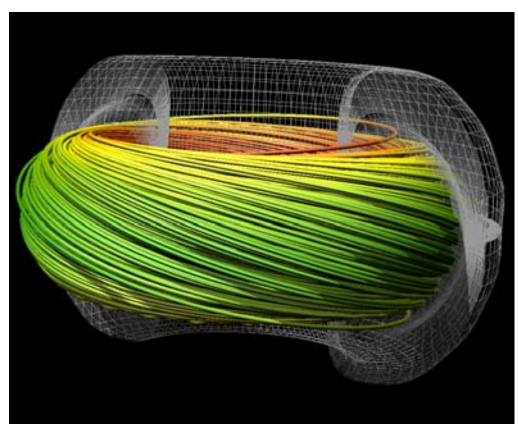

Figure 1.6. Streamlines showing the flow of the magnetic field inside a toroidal plasma chamber.

additional techniques, including Poincaré analysis (a feature-based analysis method for understanding magnetic fields in fusion simulations), time dependent streamlines, and stream surfaces.

Providing parallel data analysis support

When Ramgen Power Systems, LLC, received an allocation of time to conduct very-high-resolution computational fluid dynamics (CFD) simulations under the DOE ALCC program, Mike Matheson, a mechanical engineer and visualization liaison, was assigned to the project. Matheson's background in CFD and aerospace engineering made him a perfect fit.

The main project goal is to use the power of Jaguar to enhance and accelerate Ramgen's ability to characterize component, rotor, and stage performance and thus optimize design of its Rampressor compression technology for use in carbon capture and sequestration. However, it was obvious from the start that an ancillary project goal would have to be improving the scalability of the NUMECA FINE (flow integrated environment)/Turbo code used by Ramgen for the CFD analysis. To date, performance data have been gathered for other large-scale FINE/Turbo simulations, and a series of successful parallel scalability tests have been run as part of the development of an HPC prototype version of the software (Figure 1.7). An HPC FINE/Turbo prototype will be tested in the future, and Ramgen intends to run a significant number of Rampressor rotor and inlet guide vane interaction analyses using the prototype. Allan Grosvenor, senior engineer at Ramgen and PI for the project, says that "The insights we are gaining working with the OLCF are helping us to dramatically accelerate our research and product development." Development of a scalable HPC version of FINE/Turbo will benefit not only Ramgen, but also the broader community of designers and manufacturers that use the software.

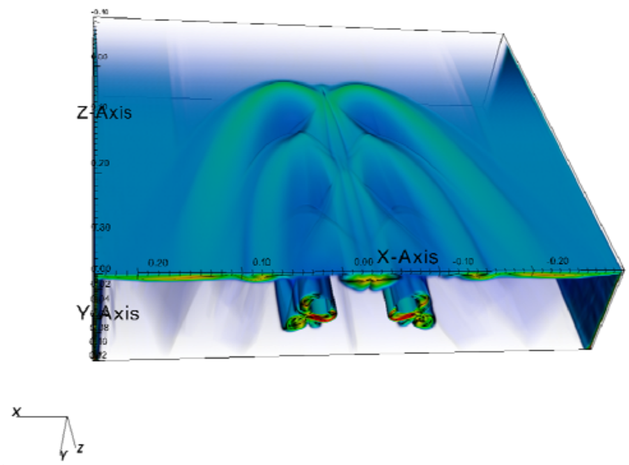

Figure 1.7. Five hundred million cell, two body simulation using FINE/Turbo. We believe this to be the largest FINE/Turbo simulation computed anywhere to date and at a high enough resolution to show strong shocks interacting with boundary layers.

Problem solving

Climate scientists are in the process of evaluating a high-resolution (T341) version of the Community Atmosphere Model (CAM) that is three times finer than typically used for climate simulation. CAM is the atmospheric component of CCSM. Using a resolution of T341, local features that are necessary to make climate predictions at the subglobal level can be resolved. When SciComp visualization liaison Jamison Daniel and computational climate scientist Kate Evans of the Computer Science and Mathematics Division used data from a T341 model run to produce a series of high-fidelity visualizations, they noticed a ringing effect emanating from regions with high values of integrated water vapor. They had not noticed this subtle effect using typical climate plotting techniques. After the two consulted with NCCS director Jim Hack, a long-time leader of climate model development and current head of Oak Ridge's Climate Science Institute, Daniel ran additional visualizations and a series of line plots to determine the source of the problem. From this analysis, they now understand this ringing to be associated with the change in time scales of the subgrid scale parameterization of cloud physics with higher resolution and have altered their solution method accordingly. The T341 modeling effort is part of a larger project on ultrahigh resolution climate simulation that Hack leads. Daniel, Evans, and Hack are currently working with collaborators at NCAR to determine the most effective way to deal with complex issues such as cloud physics within high-resolution models. A recent paper addressing this issue with lead author Mark Taylor at SNL has been submitted to the proceedings of the SciDAC 2010 conference.

1.3 COMMUNICATIONS

Communicating is one of the most important activities at the OLCF, whether it is communicating science results to the larger community or communicating tips to researchers 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:

- UAO and SciComp support services (discussed previously);
- weekly message to all users on events;
- monthly user conference calls;
- annual users meeting;
- workshops; 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).

The OLCF provides a wide range of communications products to current and potential users, the general public, and sponsoring agencies: for example, the annual report, ASCR News Roundup highlights, and so on. Below we will focus on the website, which has undergone the most significant evolution this year, and the workshops and seminar series.

1.3.1 Web Resources

UAO has deployed a dynamic new website (http://olcf.ornl.gov) to highlight the science, technology, people, and activities of the OLCF and provide enhanced access, information, and services, including system information and statistics, OLCF project details, an online newsletter, and videos. In addition, a companion support site is being developed to provide our OLCF users with allocation and account assistance, education and training modules, and a robust knowledge base.

The new site features multiple improvements, which will benefit all stakeholders and users.

- **Improved Speeds—**Improved techniques such as caching, fewer steps/queries to reach desired information, better coding, and image optimization.
- **Improved Searching—**Priority-based search, giving us the ability to configure the priority given to various parts of a post/page. Wildcards and logical operations will be supported, and results will be ordered by relevance rather than by date.
- **Simplified Navigation—**Use of techniques such as in-page tabbed navigation to decrease the total number of pages on the site, allowing us to display more information with fewer “clicks.”
- **Clearer Categories and Tags—**Clearer category structure and tags (keywords), making information easier to find.
- **Social Media—**Addition of social media tools such as Facebook to help site visitors share content with many different social networks, reaching a broader audience as well as driving new traffic to the site.

- **Custom Sidebars—**Addition of unique and completely custom sidebars featuring content related to the current page or article being viewed.
- **Featured Content—**Articles, announcements, and upcoming events that are easier to find and shown more prominently for longer durations on the OLCF homepage.
- **And More—**More features like the systems dashboard, software, related articles/highlights, article archives, links to support and knowledge base, and more multimedia are planned for the site.

We will also be able to get more meaningful information about the way the site is being used. Currently we only know how many "hits" a page gets, so we know what the users are clicking on most. However, with the new site, users can actually rate and comment on the help articles. Using their feedback, we can update articles, add additional articles, and/or delete articles that are not used. We have also assigned ownership for all articles to a UAO consultant who will be held responsible for the data contained in the articles. By assigning ownership and responsibility, the articles should stay up-to-date and be more accurate than on the previous site.

Workshops and seminar series 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. OLCF outreach to train current and future scientists and engineers is described in more detail in Section 3.3, but some of our recent activities are described below.

1.3.2 User Workshops

This year for the first time, two ORNL-based computing facilities joined the National Energy Research Scientific Computing Center (NERSC) to sponsor a Cray XT5 workshop (February 1–3, 2010, at the University of California–Berkeley). Staff from the OLCF and the University of Tennessee's National Institute for Computational Sciences (NICS) attended the Joint NERSC/OLCF/NICS Cray XT5 Workshop to train HPC and scientific communities in the use of the world's largest and most powerful leadership systems. Topics included programming effectively for the Cray

COMMUNICATION AS TRAINING

Since January 2009, the OLCF has run a science-writing internship program to optimize its outreach and communication efforts. Interns who started in July 2010 are Charli Kerns and Eric Gedenk, both seniors in journalism at the University of Tennessee (UT), Knoxville. Past interns—also UT students or recent graduates—include Caitlin Rockett (currently a communications intern for the Oak Ridge Climate Change Science Institute), Katie Freeman (currently a science news intern for ORNL's Communications and External Relations Directorate), and Beth Storey and Wes Wade, both recent writers for local newspapers. OLCF interns write about climate, astronomy, biology, physics, chemistry, and other research simulations run on leadership-class computers. To convey these scientific accomplishments to the public, they produce journalistic materials, which so far have included more than 70 news and feature articles, as well as marketing materials such as text for slideshows, boilerplates, brochures, posters, and kiosks. They also conducted research for management on such topics as social media and new users.

The OLCF science-writing internship program is managed by OLCF science writer Dawn Levy, who joined ORNL in 2007 after managing Stanford University News Service's science-writing internship program for 7 years, where she trained nearly 50 interns, most of whom went on to top media outlets including *Science, Nature, Los Angeles Times, Dallas Morning News, New Scientist,* and *Wired.* OLCF's interns are mentored by Levy and fellow science writers Leo Williams and Scott Jones. The interns' valuable contributions free these senior science writers to pursue more challenging assignments and strengthen the OLCF's overall outreach efforts. Moreover, interns trained at the OLCF are more comfortable covering HPC topics throughout their careers, have a wider net of experts to consult as sources for stories, and can parlay that education into more effective communications to the public of the importance and impact of such research.

XT5 and proper use of the new six-core CPU architecture. This marks the first time the three computing facilities have collaborated to host a workshop.

OLCF also provided the following training opportunities this past year.

- Fall 2009 Cray XT5 Workshop (December 7–9).
- Vampir Workshop (January 13–15).
- “Train the Trainer” Workshop on DDT (January 25).
- Joint Cray XT5 Workshop (February 1–3).
- Spring 2010 Cray XT5 Workshop (May 10–13).
- Visualization with VisIt 2010 (May 13).
- Crash Course in Supercomputing (June 17–18).

1.3.3 LCF Seminar Series

The LCF Seminar Series consists of monthly seminars on HPC topics, with an emphasis on petascale computing (Table 1.4). Guest speakers are typically researchers in the computer and computational sciences from outside the OLCF who interact and collaborate with OLCF staff and/or use OLCF resources. The series is a great way for ORNL researchers to interact with colleagues from around the world and a vehicle for researchers to present scientific results from use of the OLCF facilities.

Table 1.4. 2009–2010 LCF Seminar Series

Date	Series Speaker	Affiliation	Title of Presentation
August 18, 2009	Mikhail Shashkov	Los Alamos National Laboratory	Moment-of-Fluid Interface Reconstruction
October 20, 2009	Pat Teller	University of Texas–El Paso	Addressing Checkpointing and Quality of Service in HEC I/O: an Analytical Modeling and Algorithmic Approach
November 6, 2009	Victor Lotrich	University of Florida, Quantum Theory Project	ACESIII: Parallel Implementation of Coupled-Cluster Methods, a Practical Perspective
December 15, 2009	Oreste Villa	Pacific Northwest National Laboratory	Task-Based Dynamic Load Balancing and Acceleration of TCE-CCSD(T) on GPU-Enabled Systems
January 26, 2010	Markus Eisenbach	ORNL	Thermodynamics of Magnetic Systems from First Principles: WL-LSMS
February 16, 2010	Edoardo Apra	ORNL	What is a 200,000 CPUs Petaflop Computer Good For (a Theoretical Chemist Perspective)?
March 22, 2010	Ron Oldfield	Sandia National Laboratories	System Software Research for Extreme-Scale Computing
April 30, 2010	Jeff Larkin	Cray, Inc.	A Comparison of Accelerator Programming Models
May 18, 2010	William Tang	Princeton Plasma Physics Laboratory	Challenges of the Fusion Simulation Program
June 16, 2010	Stanimire Tomov	University of Tennessee	MAGMA—a New Generation of Linear Algebra Libraries for GPU and Multicore Architectures
July 1, 2010	Joel Saltz	Emory University	Toward Derivation, Management, and Analysis of Exascale Feature Sets

2. BUSINESS RESULTS

CHARGE QUESTION 2: Is the OLCF maximizing resources consistent with its mission?

Since its inception, the OLCF mission has been to deliver leadership computing for science and engineering, focus on grand-challenge science and engineering applications, procure largest-scale computer systems (beyond vendor design point), and develop high-end operational and application software in support of the DOE science mission.

During the past year, a significant upgrade in OLCF resources provided a greater than 50% increase in the resources available for OLCF users (Table 2.1). An upgrade to the Cray XT5 from AMD Opteron Barcelona quad-core processors to AMD Opteron Istanbul six-core processors was completed in November 2009. The upgrade went extremely well, and the results since have been very positive.

Table 2.1. Cray XT5 Specifications as of July 2010

System	Type	CPU Type/Speed	Nodes	Memory/Node	Node Interconnect	Cores per Node	Total CPUs	Aggregate Memory
Jaguar	Cray XT5	Opteron/2.6 GHz	18,688	16 GB	Seastar2	12	224,256	300 TB

During the upgrade, a portion of the machine remained available to users at all times. The Cray XT5 was partitioned into two, roughly equal, parts. While one partition was upgraded to the new AMD Opteron Istanbul processors, the other remained available in a production capacity for users. Only briefly during final acceptance of the collective 200 cabinets was the entire machine unavailable to users. And even then, the 263 TF Cray XT4 remained available.

Users were given access to roughly 50% of the Istanbul partition during the upgrade process to get codes transitioned as described above. The machine has been well received, and stability has proven to be surprisingly good. One contributing factor to stability is a new memory controller provided with the Istanbul processor, which has resulted in a 56% decrease in the number of memory failures ending in individual jobs aborting.

In the past year there has been a rather dramatic trend toward improved software stability. Job failures are tracked to provide feedback to the users and to watch for trends that can be investigated. Of the handful of reasons why the Cray XT5 suffered a system-wide outage, only one of them was software related. Early in the year, this software failure was a primary contributor to unscheduled downtimes, but since the bug was fixed in April, there have been no system-wide XT outages due to software. While we do see single nodes fail, most of the errors are now due to hardware, and the hardware error rate is decreasing.

Business results measure the performance of the OLCF against operational parameters. The operational metrics most relevant to OLCF business results are resource availability, resource utilization, and capability usage of the HPC resources.

To ensure that operational metrics are met or exceeded and that resources are used efficiently and effectively, the OLCF regularly measures and tunes the effects of operational policy through a series of technical and operations councils. These councils not only maximize efficiency and effectiveness, they also contribute yet another facet to customer communications and support.

Resource Utilization Council

Every successful organization has effective internal communications and collaboration—making the whole more than the sum of its parts. The Resource Utilization Council (RUC) is where all parts of the OLCF organization come together to share information and work on problems. Experience has taught us that what one part of the organization does will inevitably affect someone or something else. The RUC meets weekly, making decisions on things like Director's Discretionary (DD) awards (Section 3). It analyzes operations, including failure rates and resource utilization, with a strong user focus to help shape OLCF policies and procedures. This has led to the following service improvements and resource innovations in the past year.

- Based on reviews of resource utilization, the RUC recommended a new pullback policy to ensure computing resources continue to be used efficiently. INCITE projects that have not used a significant amount of their allocation by certain dates during the allocation year will have a percentage of the unused balance moved to a reserve pool, where it will be available for other projects. [Note: ORNL and Argonne National Laboratory (ANL) jointly agreed on the new policy, and these allocation adjustments will be made on a case-by-case basis in collaboration with the project PIs.]
- In response to analysis of user requests, the RUC recommended creation of a new external login system that enables users to log in, examine data, and compile codes even when compute partitions are down. The catalyst was analysis of Cray XT login nodes, which had not been refreshed with the latest technology, and a desire to provide more flexible, resilient login services for users. By taking the service outside the Cray XTs, the OLCF has been able to offer users better resources for code development, more reliable service, and more flexible access to both OLCF XT platforms. In consultation with OLCF HPC Operations staff and others, Cray has now produced its own version of the external login system.
- To promote leadership usage of the OLCF systems, the RUC initiated a study of queuing on OLCF systems. Empirical data in the form of queue simulations and examination of batch system logs were used to formulate a new queuing policy. Based on the results, the RUC suggested a combination policy that gives precedence to high-core-count jobs while lowering the priority of users who have more recently used the system to ensure that all projects get an equitable chance to use system allocations. The new queuing policy was implemented after the OLCF User Council reviewed it.

Operations Council

An eight-representative Operations Council ensures that day-to-day operations are safe, secure, compliant, effective, fiscally sound, and responsive to personnel and user needs. In the past year Operations Council members have provided important support and quality assurance for the daily work environment. Examples include interaction with the laboratory shift superintendent and safety services division on safety issues and providing training in critical cyber security procedures to OLCF users on proprietary and sensitive projects.

Software Council

Representatives from all OLCF groups serve on the Software Council (SWC). It grew from the desire to make the OLCF user experience as positive as possible by

- ensuring that software decisions are made in an efficient, effective, consistent manner;
- giving users a central place to go with software requests;

- ensuring that user requests are answered in an expeditious manner (1 week); and
- ensuring that new software approved for the system is promptly and efficiently loaded.

The SWC assesses user requests for new or updated versions of software to be installed on OLCF systems and ensures that all software, once loaded, is managed throughout its lifetime. Communication among SWC members is routinely carried out via e-mail, with formal council meetings once each quarter. SWC leadership rotates between representatives from the UAO, TechInt, and SciComp groups, and a subgroup ensures that software requested by users and approved by the SWC is installed correctly, with correct permissions and related issues resolved. SWTools and the software tracking system (Section 5.5) were an outgrowth of this council's mission to address user software needs, and the SWC uses both tools in conducting its work and making decisions—further contributing to the positive user environment.

User Council

The User Council is a small group—8 to 10 users—that represents the body of system users, especially with respect to issues, concerns, and suggestions for facility operation and improvements. Members are selected annually at the User Meeting in May, with officers selected biennially.

The User Council made two major suggestions this year. One was the creation of a wiki for the council to share information and discuss issues as they arise. The wiki has been created and will hopefully lead to even more collaboration among council members next year. The council also asked the OLCF to look at providing data gateways so that users could easily share their results with colleagues who are not current users and/or members of their projects. The OLCF is actively investigating a possible solution.

CHANGING THE FACE OF SCIENCE AND SHARPENING AMERICA'S COMPETITIVE EDGE

Since 2004, when DOE SC created OLCF as an open source national user facility for leadership computing, OLCF has delivered a series of increasingly powerful computer systems to the science community, beginning with a Cray X1 (3 teraflops) in 2005. Through a series of upgrades, the computing power has increased 1,000 times. This past fall, in a phased approach, the Jaguar XT5 was upgraded to 37,000 2.6 GHz six-core AMD Istanbul processors, increasing performance 70 percent over its quad-core predecessor.

The phased upgrade meant that users had access to a substantial fraction of the Jaguar XT5 platform throughout the upgrade period, minimizing impacts.

Following the upgrade, OLCF ran the benchmark program High-Performance Linpack (HPL) at a speed of 1.759 petaflops, garnering the title "fastest computer in the world." That benchmark was soon topped by an application running on Jaguar that achieved 1.84 petaflops, winning the 2009 ACM Gordon Bell Prize. Equally if not more impressive, Jaguar went on to win three gold medals in the HPC Challenge. The HPC Challenge benchmarks examine the performance of HPC architectures using kernels with more challenging memory access patterns than the HPL benchmark and thus give a better view of a system's entire performance.

The successful deployment of the Jaguar XT5 platform was not only a landmark on the OLCF road to petascale computing, but vindication of the OLCF strategy of phased scaling steps of system upgrades and new system acquisitions that allowed a concurrent phased approach to facility upgrades, system software deployment, and dramatic application scalability improvements.

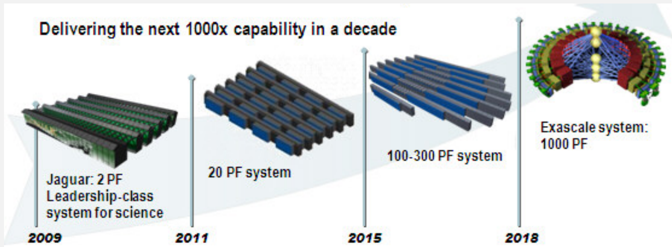

Having successfully transitioned from sustained teraflop to sustained petaflop computing, the OLCF has laid out a similar roadmap to transition to exaflop computing (figure).

The 2007–2009 chair of the User Council was Jacqueline Chen of SNL. Balint Joo of the Thomas Jefferson National Accelerator Facility was recently selected as chair of the 2009–2010 User Council. The OLCF looks forward to working with him and the rest of the council in the coming year.

2.1 RESOURCE AVAILABILITY METRICS

The OLCF tracks a set of metrics that reflect the performance requirements of DOE and the user community. These metrics assist us in monitoring system performance, tracking trends, and identifying and correcting problems at scale, all to ensure that OLCF systems meet or exceed DOE and user expectations.

- **Scheduled availability**—measures the effect of unscheduled downtimes on system availability (SA). Scheduled maintenance, dedicated testing, and other scheduled downtimes are not included in this metric. The goal is 85% scheduled availability in the first year after initial installation or a major upgrade, growing to 95% for systems in operation more than 1 year after initial installation or a major upgrade (Table 2.2).
- **Overall availability**—measures the effect of both scheduled and unscheduled downtimes on SA. The goal is 80% overall availability in the first year after initial installation or a major upgrade, growing to 90% for systems in operation more than 1 year after initial installation or a major upgrade (Table 2.3).
- **Mean time to interrupt (MTTI)**—measures scheduled service interruptions (planned maintenance or dedicated testing) plus unscheduled system interruptions from either an internal or external source.[*]
- **Mean time to failure (MTTF)—**measures time to interrupt associated with an unscheduled system interruption from either an internal or external source.[†]

In response to the 2009 OA panel's recommendation, the following two new metrics will be introduced in the FY 2011 report to track file system interrupts and failures. The scheduled and overall availability of the centralized file system should be at least as high as that of the production system; therefore, the metrics we propose for tracking the centralized file system are 95% scheduled availability and 90% overall availability.

---

[*]Mean time to interrupt (MTTI) is calculated as follows:

$$\text{MTTI} = (\text{Time}_{Total} - \text{Time}_{AO}) \div (AO + 1)\ ,$$

where

Time_{Total} = total time in period,
Time_{AO} = duration of all outages in period,
AO = total number of outages in period.

Note that **Period Start** date is set to the uptime of the most recent outage which completed within the previous month. If an outage did not complete within the previous month, the uptime of the first outage in the given period is used as the start date. **Period End** date is set to the downtime of the last outage to begin within the given period.

[†]Mean time to failure (MTTF) is calculated as follows:

$$\text{MTTF} = (\text{Time}_{Total} - \text{Time}_{UO}) \div (UO + 1)\ ,$$

where

Time_{Total} = total time in the period,
Time_{UO} = duration of all unscheduled outages in the period,
UO = total number of unscheduled outages in the period.

Period start and end dates are as noted for MTTI calculations.

- **File system interrupt—**measures any system event that causes a subset of applications to fail or a subset of data to be unavailable.
- **File system failure—**measures any system event that causes all data to be unavailable or causes all applications using the file system to fail.

Table 2.2. OLCF Computational Resources Scheduled Availability (SA) Summary 2009–2010[a]

System	CY 2009		CY 2010		
	Target SA	**Achieved SA**	**Target SA**	**Achieved SA YTD**	**Projected SA**
Cray XT5[b]	85%	95.3%	85%	94.9%	95.1%
Cray XT4	95%	97.4%	95%	97.6%	97.9%
HPSS[c]	NA	99.6%	95%	99.5%	99.6%
Spider[c,d]	NA	99.8%	95%	99.7%	99.8%

[a]Scheduled availability by calendar year (CY). CY 2010 year to date (YTD) data in Section 2 were generated from January 1, 2010, through July 31, 2010, unless otherwise noted.

[b]The Cray XT5 underwent a major upgrade in July through November 2009. CY 2009 Cray XT5 data here and in Tables 2.2, 2.3, 2.4, and 2.5 exclude this time period.

[c]A new metric to track HPSS and Spider availability was introduced this year.

[d]CY 2009 Spider data are generated from July 16, 2009, through December 31, 2009. Data prior to July 2009 were not tracked.

Table 2.3. OLCF Computational Resources Overall Availability Summary 2009–2010[a]

	CY 2009		CY 2010		
System	**Target Overall Availability**	**Achieved Overall Availability**	**Target Overall Availability**	**Achieved Overall Availability YTD**	**Projected**
Cray XT5[b]	80%	86.7%	80%	88.6%	88.6%
Cray XT4	90%	94.0%	90%	95.0%	95.4%
HPSS[c]	NA	99.3%	90%	98.3%	98.3%
Spider[c,d]	NA	96.5%	90%	98.8%	98.7%

[a]Scheduled availability by calendar year (CY). CY 2010 year to date (YTD) data in Section 2 were generated from January 1, 2010, through July 31, 2010, unless otherwise noted.

[b]The Cray XT5 underwent a major upgrade in July through November 2009. CY 2009 Cray XT5 data here and in Tables 2.2, 2.3, 2.4, and 2.5 exclude this time period.

[c]A new metric to track HPSS and Spider availability was introduced this year.

[d]CY 2009 Spider data are generated from July 16, 2009, through December 31, 2009. Data prior to July 2009 were not tracked.

The OLCF tracks MTTI and MTTF of the Cray XT4 and Cray XT5 as a means of measuring both SA and stability (Tables 2.4 and 2.5). The Cray XT4 continues to provide a stable resource for users. MTTI results for the XT5 show improvement over last year. We have experienced a slight decrease in the MTTF for the XT5 this year to date. This has primarily been due to the integration of a significant amount of hardware in late 2009, bringing it online, and getting it stabilized. As is typical after major upgrades, hardware fallout occurs as the system stabilizes and software issues appear due to increases in scale, but the system stability increases as the hardware breaks in and the software issues, such as the portals problem that was patched in April, are resolved.

Table 2.4. OLCF Mean Time to Interrupt Summary 2009–2010

System	MTTI CY 2009 (hours)	MTTI CY 2010 YTD (hours)	Change (%)
Cray XT5	42.8	48.5	+13
Cray XT4	80.2	89.9	+12
HPSS	374.8	319.9	−15
Spider[a]	NA	580.0	NA

[a]CY 2009 Spider data are generated from July 16, 2009, through December 31, 2009. Data prior to July 2009 were not tracked.

Table 2.5. OLCF Mean Time to Failure Summary 2009–2010

System	MTTF CY 2009 (hours)	MTTF CY 2010 YTD (hours)	Change (%)
Cray XT5	76.6	71.3	−7
Cray XT4	126.8	131.3	+4
HPSS	625.2	647.9	+4
Spider[a]	NA	589.4	NA

[a]CY 2009 Spider data are generated from July 16, 2009, through December 31, 2009. Data prior to July 2009 were not tracked.

Even though figures for HPSS are still impressive, the MTTI decreased this year. This was primarily due to integrating a significant amount of hardware, bringing it online, and stabilizing it.

JOULE METRIC HELPS SUSTAIN PROGRESS IN HIGH-PERFORMANCE COMPUTING

Scientific computing is critical to the breakthrough science needed to address the grand challenges of our time. Just as critical is the continuous improvement of the scientific software applications required by researchers to address ever more challenging questions. In FY 2003, the DOE SC worked directly with OMB to come to a consensus on an appropriate set of performance measures to gauge the evolution of software identified as critical to DOE mission needs. The scientific performance expectations of these requirements reach the scope of work conducted at the DOE national laboratories, and the Joule system emerged from this interaction. Ensuring compliance with these metrics, which are tracked on a quarterly basis, is an important milestone each fiscal year for the DOE ASCR Program Office as well as for the overall success of the DOE SC open science computing effort. OLCF resources—both leadership computers and expert staff—are necessary elements in the successful implementation of Joule development and performance-enhancing activities. The OLCF, through its user support model and participation in the Joule program, works hand-in-hand with researchers to produce next-generation versions of applications that can make grand challenge simulations faster and more accurate. Since implementation of the Joule metric program 5 years ago, OLCF has met all annual metric goals.

Implementation typically involves selecting four applications for testing each year. The following four applications were measured in 2009.

- VisIt (the first visualization code to undergo Joule measurement and certification)—VisIt demonstrated excellent weak scaling in each of the two visualization tasks assigned, and hardware utilization was increased 3.1 times. Additionally, two barriers to effective volume rendering were identified and addressed.
- CAM (the atmospheric component of CCSM)—Strong scaling results for CAM were outstanding: with changes to improve use of Jaguar's architecture, the software executed 2 times faster. These improvements will enable better throughput for climate scientists using the code and could have a dramatic effect on scientific productivity.
- XGC1 (a plasma physics code)—With enhancements to the software, performance results were outstanding: the software computed 4 times as many physical time steps with 4 times the number of processes in less time than previously. Improvement to the physics capability as a result of this Joule exercise is significant.
- RAPTOR (a massively parallel flow solver optimized for large eddy simulations)—As a result of performance enhancements, RAPTOR's runtime for a constant problem size was halved.

The complete FY 2009 Joule metric results were published by the OLCF in an ORNL technical report available at http://www.nccs.gov/wp-content/media/nccs_reports/FY09Q4-JouleMetric-Report.pdf.

The following codes have been selected for testing in 2010.

- Denovo—a parallel transport solver that is a first-of-a-kind, mathematically consistent, two-level approach to the multiscale challenge of nuclear reactor simulation.
- TD-SLDA—the Time-Dependent Superfluid Local Density Approximation, a recently formulated full 3-D time-dependent version of the DFT for superfluid nuclear systems.
- LS3DF—the new linearly scaling 3-D fragment code for electronic structure calculations, including simulations of solar cells.
- POP—the ocean component of CCSM.

2.2 RESOURCE UTILIZATION AND CAPABILITY USE METRICS

Allocations to center systems are made via three programs: INCITE, ALCC, and the DD program. The majority of the hours are awarded via INCITE and are granted by calendar year.

- CY 2009 Allocations: Total 598 million hours (470 million INCITE, 81 million ALCC, 47 million DD)
- CY 2010 Allocations to date: Total 1,235 million hours (950 million INCITE, 215 million ALCC,* 94 million DD)

The INCITE 2010 allocation, at 950 million hours, is 60% of the total allocated hours on the OLCF systems and represents an increase of nearly 50% over the 2009 allocation. The 2010 ALCC allocation includes time awarded to the Joule project. Through the ALCC program allocations process, 205 million hours of additional grants of time were provided to fifteen 2010 INCITE projects that had not received the full allocation originally requested.

INCITE usage for CY 2009 was about 106% of the total allocation. INCITE usage in CY 2010 to date (6/23/2010) is 57% of the allocation.

Capability (leadership) use metric

This metric measures how much of the system is used for capability (i.e., leadership computing) jobs versus smaller jobs (Figure 2.1). A single application is not considered a leadership application unless it is capable of using 20% or more of the compute cores in a single run (in year one of operation of that system or after an upgrade). Leadership usage on a system describes total consumption on the resource in terms of core-hours. The CY 2010 leadership metric is that at least 35% of the usage will be from jobs requesting 20% or more—in other words 45,000+—of the available cores.

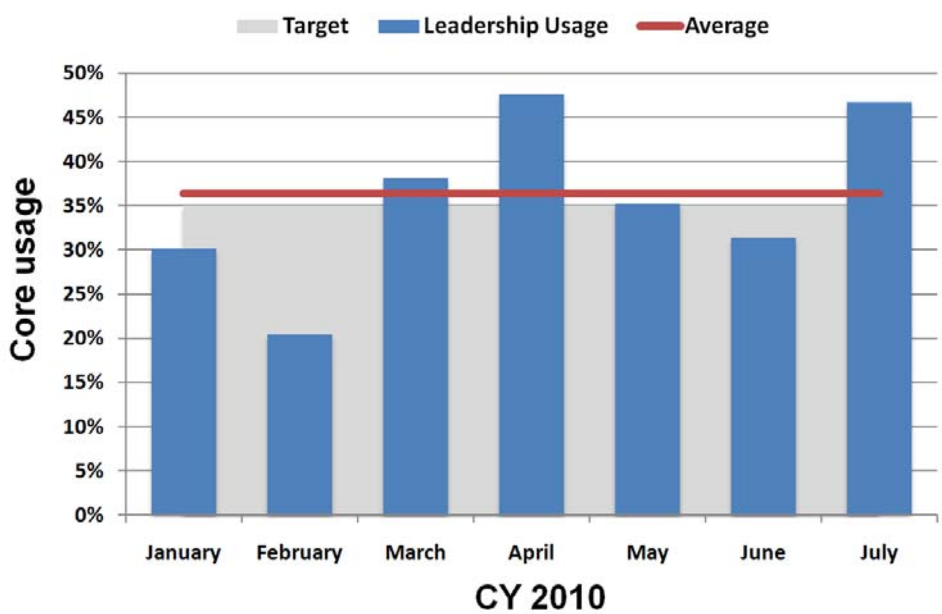

Figure 2.1. CY 2010 Cray XT5 leadership usage: Usage by jobs requesting more than 20% of the available cores.

As with other metrics, the OLCF continues to meet expectations for capability usage of its HPC resources (Table 2.6). Keys to the growth of leadership usage include the SciComp group members, who work hand-in-hand with users to port, tune, and scale code, and the ORNL support of the Joule metrics, where staff actively engage with code developers to promote application performance.

*Estimated ALCC CY 2010 allocation is based on a 12-month award period beginning in June 2010.

Table 2.6. OLCF CY 2010 Leadership Usage

Leadership Usage	Target (%)	YTD (%)
≥20% of cores	35.0	36.4

2.3 INFRASTRUCTURE

2.3.1 Networking

The first ORNL/OLCF Science Data Network (SDN) circuit was turned on in May 2009. ORNL provides the 10GE optical circuit to the ESnet Nashville hub connecting with the ESnet SDN backbone. Capacity for additional SDN circuits is reserved if needed. SDN enables dynamic provisioning of dedicated circuits between connected research facilities specifying the bandwidth and the amount of time needed for the dedicated circuit. The ORNL SDN circuit has been used for dedicated data transfer testing with ANL and ESnet. Researchers are able to move large data sets between HPC centers via the 10GE connection to ESnet. A collaboration among staff at the OLCF, NERSC, and ALCF has resulted in a 20 times increase in effective data transfer speeds over ESnet between the OLCF and NERSC.

OLCF will connect to the 100 Gb/s prototype network that ESnet staff are designing, building, and will operate. This initiative is an American Recovery and Reinvestment Act (ARRA) funded project to build an end-to-end 100 Gb/s prototype network between the DOE supercomputer centers and the metropolitan area network in New York City; it is also a network test bed facility for researchers and industry. The goal of the prototype network is to accelerate deployment of 100 Gb/s technologies and build a persistent infrastructure that will transition to the production ESnet network around 2012. This is considered a key step toward the DOE vision of a 1 TB network linking DOE supercomputing centers and experimental facilities.

Network upgrades

This past year, the OLCF deployed a high-performance network that supports the connection of OLCF infrastructure systems to diverse or redundant network switches, using all available links. This increases network resiliency, removes single points of failure, and increases effective data transfer rates by more fully utilizing available network connections.

These changes to upgrade the support networks for the infrastructure systems will increase the availability of HPC resources by minimizing the visible impact of outages and planned upgrades. The upgrades will make use of all links, increasing the capacity to servers. An increase in server utilization is coupled to infrastructure upgrades that exploit VMware so that fewer servers are deployed, and making more effective use of systems. This has an additional benefit of improving the security posture of these systems, reducing power consumption, saving floor space, and reducing cooling.

2.3.2 Storage

Storage requirements continue to grow at high rates. The current HPSS archive supports more than 12 PB of data, doubling from 6 PB a year ago. Currently the OLCF has four SL8500 tape libraries, each holding up to 10,000 cartridges, and is adding a fifth tape library in 2011, bringing the total storage capacity up to 50,000 cartridges. The libraries house a total of 24 T10K-A tape drives (500 GB cartridges, uncompressed). The tape drives can achieve throughput of 120–160 MB/s.

An average of 17 TB is written to tape every day, and this figure continues to grow exponentially. As storage, network, and computing technologies change, the OLCF's storage system is evolving to take advantage of new equipment that is both more capable and more cost-effective. The OLCF is actively examining the topics of media refresh, data retention policies, and archive system performance. We are actively engaged in a significant HPSS upgrade (Section 5.6).

2.4 ENERGY SAVINGS FRONT AND CENTER

Maximizing resources involves more than just providing computer power when and where it is needed. It also means being good stewards of the resources we have. The supercomputers at ORNL's computing complex, which include the world's fastest academic supercomputer, the University of Tennessee's Kraken, and the world's largest dedicated resource for climate prediction, all consume substantial amounts of energy with equally large demands for robust cooling and support infrastructure.

The Computational Sciences Building (CSB) housing Jaguar and Kraken was among the first computing facilities in the country to be LEED (Leadership in Energy and Environmental Design) certified. As a result of careful engineering practices, the CSB has a power usage effectiveness (PUE) of about 1.25, compared to an average of about 1.8 among other large-scale data centers. In practical terms, this means that within the facility each 1 MW used to power the machines, requires just 0.25 MW for supporting functions including the removal of waste heat, lighting, and other ancillary facility services. ORNL has a second computing center that was built shortly after the CSB. This facility adheres to even more rigorous engineering practices and is LEED-Gold certified.

Since completion of the facility in 2004, the OLCF continues to identify ways to reduce its resource footprint even more, harnessing energy savings wherever possible.

The cooling method used for the largest systems is a principle contributor to the low PUE. Cray's ECOphlex cooling system, used for the three largest computer systems in the facility, transfers waste heat directly to chilled water using a method that is conservatively 10 times more efficient than traditional systems that rely solely on air for temperature control.

ORNL also manages the mechanical systems very carefully. One example of this is the careful evaluation of increasing the inlet water temperature to the ECOphlex system. The Cray systems that use ECOphlex are very significant consumers of chilled water. The delivery of warmer water, and the ability to maintain or even increase the temperature differential across that cooling system, allows ORNL to recognize considerable cost savings in the Central Energy Plant without any increased risk to the Cray systems.

Mechanical system improvements continue to yield good savings. We recently completed the replacement of the high volume pumps and cooling fans on the original building chillers so that we could deliver a system-wide variable primary pumping scheme. The savings from this one activity are estimated at up to $100,000 per year, creating an opportunity for a rapid return on investment and sustained savings for the life of the facility. Other CSB computer room improvements, many of which addressed small mechanical system improvements such as the widespread use of VFDs on the computer room's air conditioning units and other equipment, are calculated to save another $355,000 per year compared to the original building design.

ORNL has moved aggressively during the short life of the facility to identify areas where we could increase operational efficiency. These measures have been taken to ensure the OLCF supercomputing program makes the most efficient, effective use of resources and taxpayer dollars. We are determined to get the most computing bang for the power buck among the leading HPCFs, allowing us to tackle big science, not only more quickly, but also more efficiently.

We will use the experience we have gained building and managing current facilities as we build our next center to support exascale systems.

3. STRATEGIC RESULTS

CHARGE QUESTION 3: Is the OLCF meeting Department of Energy Strategic Goals 3.1[*] and 3.2[†]?

DOE has an ambitious strategic plan comprising five strategic themes and sixteen strategic goals designed to help the agency successfully achieve its mission and vision. Strategic theme 3, "Scientific Discovery and Innovation," is aimed at strengthening U.S. scientific discovery and economic competitiveness and improving quality of life through innovations in science and technology. Its two strategic goals, "Scientific Breakthroughs" and "Foundations of Science," include objectives emphasizing

- scientific discovery and revolutionary approaches to the nation's grand challenges in areas like national security, energy, and environmental quality;
- delivery of scientific facilities, capabilities, and infrastructure required to maintain U.S. scientific primacy; and
- training future generations of scientists and engineers.

A Science-Based Case for Large-Scale Simulation (SCaLeS), the report produced by DOE SC in 2003,[‡] looked forward to a petascale era where simulation would truly begin to become a "peer methodology" with experiment and theory for scientific inquiry. The past year witnessed the dawning of the petascale era at the OLCF as the peak performance of Jaguar grew from 1.64 petaflops to more than 2 petaflops. The physical fidelity of many simulations performed at the OLCF in the past year is astonishing when compared to the SCaLeS predictions of a few short years ago. Detailed, predictive simulation science—the articulated hope for the petascale in the SCaLeS report—is, indeed, starting to be realized by many scientific disciplines. In this section, we will describe and, in some measure, quantify this promise by exhibiting how the OLCF met Strategic Goals 3.1 and 3.2.

3.1 SCIENCE TRACKING

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. OLCF Joule metric activities are described in Section 2. The facility also collects quarterly reports from users.

---

[*] Strategic Goal 3.1, Scientific Breakthroughs: Achieve the major scientific discoveries that will drive U.S. competitiveness; inspire America; and revolutionize approaches to the nation's energy, national security, and environmental quality challenges.

[†] Strategic Goal 3.2, Foundations of Science: Deliver the scientific facilities, train the next generation of scientists and engineers, and provide the laboratory capabilities and infrastructure required for U.S. scientific primacy. (DOE's strategic plan, including both Strategic Goal 3.1 and Strategic Goal 3.2, is available at http://www.cfo.doe.gov/strategicplan/strategicplan.htm.)

[‡] U.S. Department of Energy, Office of Science, *A Science-Based Case for Large Scale Simulation*, 2 vols., 2003/2004 (volume 1 available at http://www.er.doe.gov/ascr/ProgramDocuments/Archive/SCaLeSReportVol1.pdf; volume 2 at http://www.er.doe.gov/ascr/ProgramDocuments/Archive/SCaLeSReportVol2.pdf).

[§] Panel recommendations can be found in the full report of the committee, *Advanced Scientific Computing Advisory Committee Petascale Metrics Report*, 28 February 2007, available at http://www.er.doe.gov/ascr/ascac/Reports/PetascaleMetricsReport.pdf.

Table 3.1. Publications by Calendar Year

	2008	2009
Number of refereed publications based on the use (at least in part) of OLCF resources	369	401

In addition to works published by our users, the OLCF has a varied and active publication activity designed to disseminate information to as wide an audience as possible. In the past year the center has produced several reports, including the 2009 OLCF Annual Report, *FY 2009 Annual Report of Joule Software Metric, Science at the Petascale 2009*, and *Preparing for Exascale: ORNL Leadership Computing Facility Application Requirements and Strategy*. The reports can be found at http://www.nccs.gov/media-center/nccs-reports/.

2009 Gordon Bell Prize team sets the stage for magnetic-materials design

Calculating materials properties using first principles has become a standard tool over the last few decades. One shortcoming of these density functional calculations, however, is that they only describe material behavior at absolute zero (about −460°F)—not very practical for most applications. The OLCF's Markus Eisenbach is part of an international team studying the behavior of magnetic systems at finite (i.e., real-world) temperatures such as the Curie temperature—the temperature at which materials lose their magnetism. The team developed the method to do this by combining a classical density functional method (LSMS) with a classical Monte Carlo method known as Wang-Landau (WL). Prior to this, working at finite temperatures was possible only with very simple systems, a severe limitation. Eisenbach says that "These first principles calculations are orders of magnitude more computationally demanding than previous models," and that "it is only with a petascale system such as Jaguar that calculations like this become feasible." The method, known as WL-LSMS, used more than 223,000 of Jaguar's 224,000-plus available processing cores and achieved 1.84 petaflops (80 percent of Jaguar's rated peak performance of 2.33 petaflops), winning the 2009 ACM Gordon Bell Prize. More importantly, by accurately revealing the magnetic properties of specific materials, the project promises to boost the search for stronger, more stable magnets, thereby contributing to advances in areas such as magnetic storage and the development of lighter, stronger motors for electric vehicles. Other areas that may benefit from this research include the design of lighter, more resilient steel and the development of future refrigerators using magnetic cooling.

3.2 SCIENTIFIC ACCOMPLISHMENTS

From its earliest days as a production center, the OLCF has approached the delivery of science on its computational resources as a collaborative enterprise. Projects are provided essential support ranging from the immediate concerns of system configuration and usage to advanced algorithmic, performance, and numerical implementation support from SciComp liaisons. This total support model, pioneered at the OLCF, has been widely recognized as a best practice for current and future centers. See Section 1.2 for examples of many success stories, ranging from optimizing code performance and scalability to assisting with code development and algorithms, culminating in faster, more effective use of leadership systems to facilitate scientific problem-solving.

Computational scientists and other experts at the OLCF have engaged researchers worldwide to address the leading challenges facing the nation, and the scientific results stemming from this collaborative effort show that the OLCF strategy is paying off. We are confronting and answering big science questions and grand challenges—in energy, climate, materials science, physics, chemistry, and environmental science. Some of these accomplishments are highlighted below, including scientific results coming from applications with delivered performance well in excess of a petaflop.

3.2.1 Energy

Nuclear engineering for energy independence

In spite of the fact that conventional nuclear reactors (i.e., those based on nuclear fission) have been around for more than 60 years, the public has had difficulty "warming" to the technology. Accidents such as those at Three Mile Island in the United States, Windscale in Great Britain, and Chernobyl in the former Soviet Union have contributed to public concerns about nuclear power. Over the years, efforts to improve reactor design, minimize the need for waste storage (another source of concern), and increase safety have been hampered by the expense of experiments in this area and the enormous complexity of supercomputer modeling—seven independent variables leading to billions of spatial elements, hundreds of angles, and thousands of energy groups. These complexities have meant that, until recently, simulations were limited to approximations and used averaging methods known as homogenization, which limited their ability to explain localized behavior in reactors and thus their usefulness.

With the advent of petascale machines such as Jaguar, researchers such as ANL's Dinesh Kaushik, an Early Science award recipient, will be able to move toward more detailed, realistic simulations of reactors. Kaushik and his colleagues are using Jaguar to increase our understanding of reactor processes using the UNIC code, developed under DOE's Advanced Modeling and Simulation program. The code will provide progressively more detailed descriptions of nuclear reactor core and associated processes (Figure 3.1). The advance will benefit the design not only of traditional light-water reactors, but also fast reactors, which are likely to use existing reactor waste for fuel, thus converting and reducing the inventory of existing spent reactor fuel. To date the UNIC code has run on 131,072 of Jaguar's 224,000 processor cores for two reactor problems. Kaushik says allocation on Jaguar "will allow us to carry out more realistic reactor simulations, resulting in less uncertainty in the crucial reactor design and operational parameters."

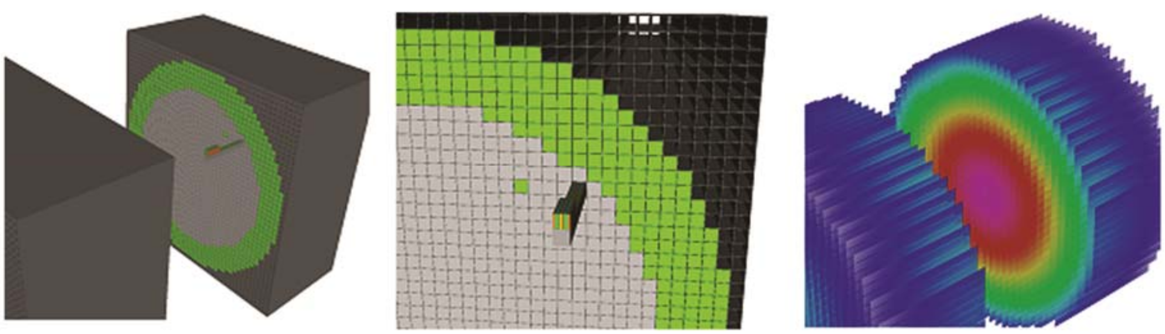

Figure 3.1. Two pictures (left and center) of ZPR 6/6A geometry and uranium-235 plate power distribution (with separated matrix halves). The gray indicates the matrix tube and drawer fronts that are loaded into each tube position. The solid green squares are 2-inch depleted uranium metal blocks directly loaded into the tubes surrounding the main core and acting as a neutron blanket. The plot at the right shows the enriched uranium plate power with the matrix halves separated. *(Images courtesy of Dinesh Kaushik, ANL.)*

The future of nuclear energy: harnessing the power of the stars

Recent fusion simulations on Jaguar using the XGC1 particle-in-cell code and an INCITE award of time have verified what has long been speculated: The temperature and turbulence at the edge of a fusion plasma affect the temperature and turbulence of the plasma core (Figure 3.2). This finding has enormous implications in the quest for fusion energy. Plasmas are susceptible to all sorts of temperature and density (turbulence) fluctuations. These fluctuations, if too strong, can easily degrade the plasma confinement, thus weakening the fusion reaction and with it the hopes of economically feasible energy production. Understanding the dynamics between the temperature and turbulence in the edge and their influence on the core will be key in eventually generating commercially viable fusion power.

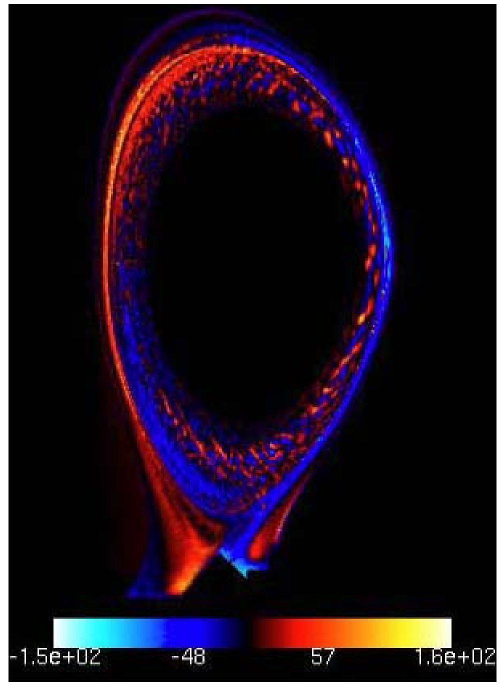

Figure 3.2. Ion temperature gradient turbulence on a tokamak fusion reactor. *(Image courtesy of Scott Klasky, OLCF.)*

The latest simulations verified that turbulence in a well-confined edge can penetrate the core and boost its temperature, something which had long been postulated. Given that a major problem in fusion is maintaining the core's temperature (ten times hotter than the surface of the sun) while keeping the plasma edge that is in contact with the wall cool, the simulation delivered good news. Generally the hotter the core, the better, and the fact that edge turbulence carries the high-temperature property to the core is a plus for maintaining the fusion reaction. A better understanding of the profile of both the edge and core is necessary if the upcoming prototype ITER fusion reactor is to function optimally. The current simulations used 20,000 of Jaguar's cores. In the future, the project goal is to simulate the entire ITER device using all Jaguar's cores. Visualizing the workings of the whole device, revealing the numerous relationships at play in the complexities of a working fusion reaction, will be a major contributor to ITER's success. Because of the complexities of all such fusion simulations, only platforms like those at the OLCF are capable of handling them in a reasonable timeframe. Project lead C. S. Chang of New York University has enthused that "The purpose of the code and the purpose of the machine fit perfectly together," . . . and for simulations of this type "Jaguar is number one."

Solar energy: harnessing the power of our star

One reason solar energy use has not become more widespread is the expense and efficiency of current solar cells. Finding better materials from which to make solar panels is a big part of the battle. Scientists from LBNL are using Jaguar to better understand solar panel materials at the atomic level. Led by Lin-Wang Wang, an Early Science award recipient, the research team is using Jaguar to model the nanostructures of potential materials.

One material of particular interest is zinc tellurium oxide, a semiconductor alloy. The team is exploring the role of oxygen atoms in the material, specifically whether the oxygen atoms will introduce an intermediate electron state in the middle of the zinc tellurium band gap (Figure 3.3). Theoretically such a system could increase solar cell efficiency dramatically.

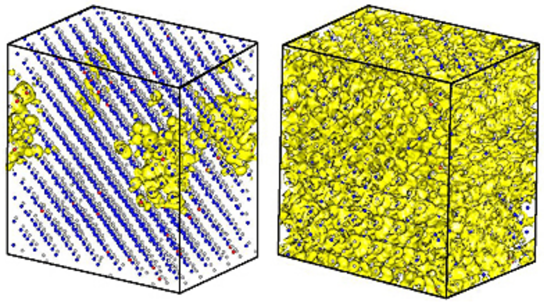

Figure 3.3. Intermediate electron states within the zinc tellurium band gap introduced by the addition of oxygen atoms to the alloy. The figure on the right shows electrons in the material that are available to conduct electricity.

"Combining new algorithms with big computers like Jaguar, we will eventually be able to simulate how the electron moves in a nanosystem, from its excitation after the sunlight absorption, to its transport to the surface and interface, and for some of them being trapped in defect states, and some of them being collected by the electrodes to generate electricity," Wang said. "Such detailed understanding is essential for designing new solar cells. It has taken 30 years for people to understand fully the simple thin film silicon solar cell. Hopefully, with the large-scale simulation, it will take less time to understand the more complicated nanocells."

3.2.2 Earth Sciences

Exploring the limits of climate prediction

In 1969 scientists from the Geophysical Fluid Dynamics Laboratory (GFDL), a branch of NOAA in Princeton, New Jersey, published results from the world's first climate model. Now, 40 years later, another Princeton/GFDL team is using a descendant of that original model and the power of OLCF's Cray XT5, Jaguar, to simulate and assess both natural and anthropogenic causes of climate change at exceptional resolutions—50 km and higher—high enough for a truly local view of weather and climate (Figure 3.4).

The "model" the team is using is actually two models built on a flexible framework that allows different components of the climate system to be modeled by multiple scientists and code developers and assembled (i.e., coupled) in a variety of ways. The models run independently but concurrently, with data exchanged about every 2 hours (compared to a 24-hour timeframe for most such models).

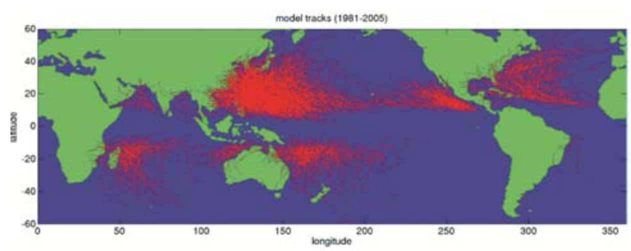

(a)

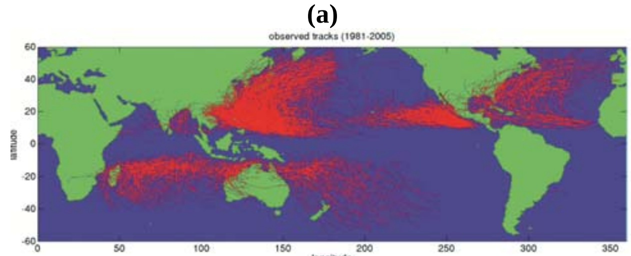

(b)

Figure 3.4. Simulation of (a) global hurricane climatology and response to global warming at 50 km resolution compared to (b) actual data for the same period. *(Images courtesy of M. Zhao, I. Held, S.-J. Lin, and G. Vecchi, NOAA GFDL.)*

Coupling climate models is computationally expensive, but the Jaguar XT5 gave the team, led by Venkatramani Balaji, the power and speed it needed to frequently link its climate models.

The project, dubbed CHiMES (for Coupled High-Resolution Modeling of the Earth System), is a collaborative effort between DOE and NOAA. Using an INCITE program allotment, in 2009 the ChiMES team ran about 500 years' worth of coupled-model simulations on Jaguar. Scaling its high-resolution models from 60,000 to 100,000 cores, the team was able to realistically duplicate the statistical behavior of hurricanes, accurately simulating their seasonal peak in September. In 2010 the team plans to use an additional allocation to determine whether decadal predictability is possible and to bring even higher-resolution models online—ones capable of resolving fine-scale weather events, such as extreme hydrological events, on the scale of individual cloud systems.

Earthquake simulation rocks southern California

California takes earthquakes very seriously. The state straddles two major tectonic plates and is subject to relatively frequent, often major, potentially devastating quakes.

It should come as no surprise, then, that the most advanced simulation of an earthquake ever performed on a supercomputer focuses on California and its San Andreas Fault. A team led by Southern California Earthquake Center (SCEC) director Thomas Jordan is using an INCITE award on Jaguar to simulate a 6-minute, magnitude-8 quake shaking a 125,000-square-mile area of southern California to assess its impact on the region, which includes 20 million people. Magnitude 8—50% greater than the temblor that destroyed San Francisco in 1906 and 30 times more powerful than the quake that devastated Haiti in January—was selected because it is one of the largest quakes that could plausibly hit Southern California.

Jaguar was required for the simulation for two reasons: (1) the size of the region being studied, which the simulation divided into 435 billion 40-cubic-meter cells, and (2) the frequency of the seismic waves, which the simulation was able to calculate up to 2 Hz—or 2 cycles per second—without resorting to approximation. Previously, no earthquake simulation of this scale has been able to directly calculate earthquake waves above 1 Hz. According to San Diego State University computational scientist Yifeng Cui, each doubling in wave frequency requires a 16-fold increase in computational resources. However, building engineers use waves up to 10 Hz in their analyses, so this is just one step along the path toward the larger goal of similar simulations at even higher frequencies.

The project conducted its first Jaguar simulation in April, running for 24 hours and taking advantage of nearly all of Jaguar’s 224,000-plus processing cores. The simulation reached 220 trillion calculations per second, or 220 teraflops, more than twice the speed of any previous seismic simulation.

In time, these simulations will contribute significantly to the information used by the state’s building designers and emergency agencies to prepare for future earthquakes, and this knowledge will ultimately be useful to scientists and other experts looking at earthquake-prone regions across the globe, not just in California.

Validating the tools for studying climate change

A multinational interdisciplinary team is attempting the world’s first continuous simulation of 21,000 years of Earth’s climate history—from the last glacial maximum to the present—using a state-of-the-art climate model known as CCSM. The model, a global climate model that includes coupled interactions between atmosphere, oceans, land, and sea ice, was developed with funding from the National Science Foundation, DOE, and the National Aeronautics and Space Administration. The group also plans to extend the simulation 200 years into the future to forecast climate change (Figure 3.5).

Most climate simulations so far are discontinuous, amounting to snapshots of century-sized time slices taken every thousand years or so. Such simulations are incapable of simulating abrupt transitions occurring on centennial or millennial timescales. INCITE awardees and project leads Zhengyu Liu of the University of Wisconsin–Madison and Bette Otto-Bliesner of NCAR are using Jaguar to, in effect, stitch together a continuous stream of global climate snapshots and recover the virtual history of global climate. According to Liu, this is the “most serious validation test of our model capability for simulating large, abrupt climate changes, and this validation is critical for us to assess the model’s projection of abrupt

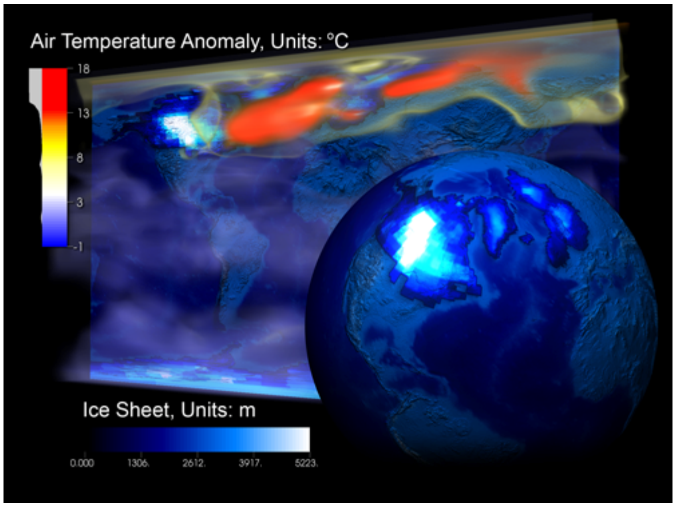

Figure 3.5. Simulations showing deglaciation during the Bølling-Allerød, Earth's most recent period of natural global warming. *(Image courtesy of Jamison Daniel, OLCF.)*

changes in the future." More accurately depicting the past means clearer insights into the climate outlook for the future. "Our simulation is an important step in assessing the likelihood of predicted abrupt climate changes in the future because it provides a rigorous test of our model against the major abrupt changes observed in the recent past," Liu says. The project was described in an article in *Science* last year.*

3.2.3 Physical Sciences

Fundamental nuclear physics: understanding nature's premier clock

Carbon-14 has been invaluable to the study of man's prehistory—both because it exists in all living things (and therefore things that were once living) and because of its long half-life, much longer than other light elements. Now, through an Early Science award and an INCITE project, an ORNL led team is using Jaguar to examine the carbon-14 nucleus. The team, which includes the OLCF's Hai Ah Nam, James Vary and Pieter Maris of Iowa State University, and Petr Navratil and Erich Ormand of Lawrence Livermore National Laboratory (LLNL), hopes to explain carbon-14's long half-life—about 5,700 years—and advance our understanding of what holds all nuclei together (Figure 3.6). The theoretical models people have been using to describe other light nuclei don't do as well for carbon-14, which means the models aren't explaining all of the physics.

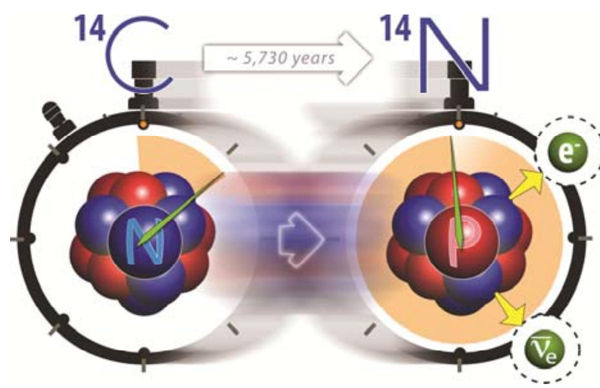

Figure 3.6. Decay of carbon-14 into the stable, nonradioactive isotope nitrogen-14.

Nam and her colleagues have used Jaguar to dissect the secrets of carbon-14 with a nuclear shell model application known as Many Fermion Dynamics, nuclear (MFDn), created by Vary at Iowa State. According to Nam, MFDn is an especially good code for this application because it scales very well. The team is using nearly 150,000 computing cores on the project, and the application is ready to scale to even more cores as they become available.

Jaguar's unprecedented power allows the team to depart from other nuclear structure studies in a variety of respects, including the incorporation of three-body forces. "Previously we could only consider two-nucleon interactions because the number of combinations needed to describe all the different interactions is really big, even for only two particles at a time," Nam explained. "And while two-particle interactions are the dominant way that these particles interact, there are some nuclear phenomena, like the half-life of carbon-14, that can't be explained using a two-nucleon interaction only. Three-particle interactions or

*Z. Liu, et al., "Transient simulation of last deglaciation with a new mechanism for Bølling-Allerød warming," *Science,* 325(5938), pp. 310–314 (2009).

higher can also be at play." Jaguar makes these calculations possible not only because of its speed but also because, at 300 TB, it has the memory required for the calculations. "These types of calculations for carbon-14 were previously not possible because it's a memory-intensive calculation," explained Nam. "Accounting for the three-nucleon force amounts to storing tens of trillions of elements . . . that's hundreds of terabytes of information." By making use of Jaguar's power, the team hopes to push us a little closer to an understanding of the atom's nucleus—crucial for answering many fundamental questions such as the origin of elements in the universe.

Water, water everywhere . . .

A team led by the OLCF's Edoardo Aprà and composed of members from ORNL, Pacific Northwest National Laboratory (PNNL), Australian National University, and Cray Inc. explored the bonding properties of water molecules, earning a position as finalists in the 2009 Association for Computing Machinery (ACM) Gordon Bell Prize competition, which honors the world's highest performing scientific computing applications. The team used a code created in the 1990s by PNNL for solving environmental restoration problems. Using the code, NWChem, the researchers were able to discern the molecule's lowest energy configuration—its most stable one. Aprà's simulation of a 24-molecule cluster is the first to explore these bonds by means of the first-principles quantum chemistry technique known as coupled cluster.

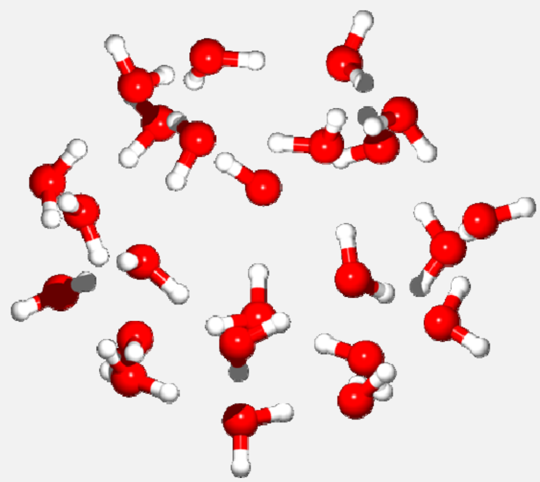

A 2009 Gordon Bell Prize finalist used the Cray XT5 to accurately study the electronic structure of water.

The unprecedented power of the Jaguar system was necessary for these calculations because the weak bond between water molecules cannot be accurately described by other, less demanding, computational approaches. The simulation required 150 TB of memory for a sustained performance of 1.39 petaflops on 223,200 processors and ran for a wall-clock time of about 3 hours.

Water plays an essential role in several key chemical and biological processes, so accurate models like this one are crucial to understanding, controlling, and predicting the physical and chemical properties of complex aqueous systems. The team will make its results available to other researchers, who will be able to use these highly accurate data as inputs to their own simulations.

3.2.4 Life Sciences

The human body comprises trillions of cells, and over the course of a life time they replicate countless times, with great fidelity most of the time. However, surprisingly little is known about this process. Now, a team of researchers led by Ivaylo Ivanov of Georgia State University is using their INCITE award on Jaguar to unlock the secrets of cell replication, modification, and repair. Their findings, which they shared this spring in an article in the *Journal of the American Chemical Society*,[*] reveal that a ring-shaped protein known as a proliferating cell nuclear antigen, or "sliding clamp" to biologists, plays a pivotal role in the process (Figure 3.7).

[*]J. A. Tainer, A. McCammon, and I. Ivanov, "Recognition of the Ring-Opened State of Proliferating Cell Nuclear Antigen by Replication Factor C Promotes Eukaryotic Clamp-Loading," *JACS*, 132, pp. 7372–7378 (2010).

"This research has direct bearing on understanding the molecular basis of genetic integrity and the loss of this integrity in cancer and degenerative diseases," says Ivanov, whose investigation was supported by the Howard Hughes Medical Institute and the National Science Foundation's Center for Theoretical Biological Physics.

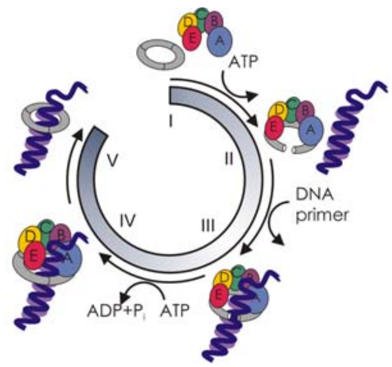

Figure 3.7. The clamp loading cycle. (*Image courtesy of I. Ivanov, Georgia State University.)*

"Sliding clamps and clamp loaders are part of the replisome—the molecular machinery responsible for the faithful duplication of the genetic material during cell division," explains Ivanov. "The replisome is very complex and dynamic, with interchanging parts. It's an incredibly challenging system to understand." Simulating just a few of its constituent parts—the clamp/clamp loader assembly—required a system of more than 300,000 atoms. "To make progress simulating the system in a reasonable amount of time, we needed access to large-scale computing." Hence the need for Jaguar's petascale power.

An improved understanding of the replisome may make it possible to exploit differences among organisms as diverse as viruses, bacteria, plants, and animals. Although clamp loaders from the different kingdoms of life share many architectural features, significant mechanistic differences exist between the various clamp-loading machines, specifically in the ways ATP is used. Drugs targeted to the clamp loader could selectively inhibit replication of viral DNA in diseases such as chickenpox, herpes, and AIDS without interfering with DNA replication in normal human cells. Similarly, in processes with increased DNA replication, such as cancer, inhibiting clamp loading might produce therapeutic effects without unwanted side effects.

3.3 TRAINING FUTURE GENERATIONS

Scientific and technological accomplishments on leadership computers would be impracticable without expert support for researchers and training—both for current and the next-generation user. In Section 1 we described the UAO and liaison support models that promote effective use of the computer system for cutting-edge science. The OLCF also provides training, from the graduate to the advanced research level, to users and potential users of leadership systems. The center uses an integrated approach that includes workshops, internships, tutorials, and educational outreach to train students and users in high-performance computational science. Eighteen workshops and seminars were held or sponsored by the OLCF in FY 2010 (Sections 1.3.2 and 1.3.3). These included the following.

- **Three OLCF-organized Cray XT5 workshops (in December, February, and May)**. OLCF, NICS, and vendor staff worked with both current and prospective users to develop and hone their XT5 skills and to introduce new, petascale-enabling software tools.
- **Visualization with VisIt (May 13)**. Attendees used their own simulation data as they worked with OLCF visualization experts to learn this breakthrough parallel visualization tool.
- **Crash Course in Supercomputing (June 17–18)**. OLCF staff taught beginning and advanced courses that introduced students to the basics of programming in the Unix environment and to the skills required for programming in a parallel environment.

The OLCF maintains a broad program of collaborations, internships, and fellowships for young researchers. For 2009–2010, the OLCF supported more than 30 faculty, student interns, and postdoctoral researchers, providing a nurturing environment where they could gain valuable work experience while working with leaders in the field. OLCF interns and postdoctoral employees have contributed in a tangible way to OLCF projects and objectives, further demonstrating the quality of the learning environment provided.

OLCF user assistance analyst Robert Whitten taught an HPC survey course by videoconference to students at Morehouse College, a historically black college (HBCU) in Atlanta; other participants included Clark Atlanta University and Texas Southern University, also HBCUs. The course focused on HPC as applied to scientific computing and discussed accessing HPC resources, as well as developing and executing applications for those resources. Course lectures and presentations gave students a foundation for pursuing more advanced learning in HPC. During the course, guest lecturers helped students see the HPC field from a variety of viewpoints.

Over the years, the OLCF has been a strong supporter of ASCR's Research Alliance in Math and Science (RAMS) program. The program is designed to provide collaborative research experiences between ORNL researchers and underrepresented students majoring in computer science, the computational sciences, math, and related disciplines. The RAMS program seeks to improve U.S. competitiveness in research and to increase the number of U.S. citizens from underrepresented populations who hold advanced degrees in science, mathematics, engineering, and technology.

Each student is assigned an ORNL research scientist mentor and conducts a project of mutual interest, often an important aspect of a larger ongoing project, over the 12-week internship period. Students keep journals of their activities and experiences, attend seminars and workshops, summarize the results of their research projects in papers and oral presentations, and participate in poster presentations at ORNL and national conferences. The program is administered through the Computing and Computational Sciences Directorate office and has provided research and conference participation

Earn, learn, return . . .

"Awesome," "amazing," "wonderful," "exciting" . . . these are just some of the superlatives used by OLCF staff members and student intern Yashema Mack to describe each other.

Mack, a second-year intern and master's candidate from Florida A&M University, first became aware of ORNL and OLCF at the 2008 Tapia Celebration of Diversity in Computing conference. Because of what she heard there, Mack applied to the Research Alliance in Math and Science (RAMS) program the same day and was accepted in short order—receiving the last open slot for 2009.

Working with her OLCF mentor, Rebecca Hartman-Baker, Mack created a web interface for the software tracking system (Section 5)—from scratch. Mack said "It was an exciting project because I had never done anything like it before and never worked with PHP [scripting language] before." The RAMS program also gave her the opportunity to build a SharePoint site—again, something she hadn't done before and for which she received a "superlative" rating in a vote of her peers. Echoing Mack's enthusiasm, Hartman-Baker said that what Mack accomplished in one summer had been expected to take at least 6 months.

Mack said that the best part of the intern experience was "Working with NCCS staff . . . picking their brains. . . . They are awesome, fun, funny . . . great motivators. . . . You don't find this kind of work atmosphere anywhere." She added that there's a hidden benefit to interning at ORNL: ORNL name recognition can give a "boost" to résumés. She believes the many unsolicited calls she has received from major corporations are the result of interning here. "ORNL is known worldwide, and people know what it means to have worked here."

When asked whether she would recommend interning here to others, Mack's answer was an unequivocal "YES!" "There's nothing like hands-on training," she added, "and this is one of the best places to get it."

Bearing this out, Mack is back this summer, by mutual request—under ORNL's Nuclear Engineering Science Laboratory Synthesis program—working again with Hartman-Baker, this time on the Joule metric program (Section 2).

opportunities to students from both majority- and minority-serving institutions across the United States, including Puerto Rico. In the summer of 2009, OLCF staff mentored several RAMS students.

In addition, the OLCF/ORNL participates in a number of formal programs designed specifically to foster and develop the next generation of scientists.

For example, the OLCF partners with the Appalachian Regional Commission (ARC), a federal–state partnership whose goal is to support the educational development of the Appalachian region. The OLCF–ARC partnership seeks to open up college and career options to underprivileged and minority high school students. Through the program, the OLCF accepts high school students and college and high school instructors annually from an area encompassing 15 states for a 2-week workshop, during which the students and instructors participate in activities and projects in areas such as biology, astronomy, and computational sciences. As one of last year's participants remarked, the program "filled in the gaps with my programming" and helped bridge "the gap between high school and college education."

The OLCF also participates in the Jason project, National Geographic's award-winning middle school science education program. Middle school students from the Jason Project visited the OLCF and the visualization center when they came to ORNL to work with Virginia Dale, director of the Center for Bioenergy Sustainability, and her team. Students learned how Jaguar is advancing the study of the combined economic and environmental impacts of growing switchgrass, a potentially important renewable energy crop. Students also viewed detailed National Aeronautics and Space Administration satellite images on Everest that demonstrated seasonal vegetation dynamics across major geographic features like the Mississippi Valley.

Feedback such as the following from one of this year's Jason Project participants tells us we are on the right track in educating the engineers and scientists of the future.

> *It's incredible to think that we were standing in a place where such cutting-edge scientific discoveries are being made. While the relationships between different areas of science may not always be obvious, our research at Mount St. Helens and in Tennessee with Oak Ridge's supercomputers has helped us to make connections between weathering, erosion, agriculture, soil, biofuel, and climate change. It was a truly inspiring experience.*—Karina Jougla, a National Student Argonaut, Carpinteria, California.

In 2010, we once again conducted a distance education course on HPC, designed as a survey course to provide an introduction to the concepts, tools, and methods for using high-performance computers to solve scientific problems. Three colleges and universities, and a total of 30 students, participated. This was just one of a range of courses and training opportunities offered over the past year and discussed in the following sections.

3.4 DELIVER THE SCIENTIFIC FACILITIES

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 Innovative and Novel Computational Impact on Theory and Experiment

The INCITE program is in its seventh year of operations. In 2010 the program was transitioned to the Argonne and Oak Ridge Leadership Computing Facilities, and it is managed by Julia C. White of ORNL. In CY 2010, 950 million processor hours were awarded on the Jaguar resource to 45 projects.

The 2010 INCITE peer-review was carried out by more than 80 researchers worldwide—national fellows, department chairs, and senior-level managers, all leaders in their fields. Whereas in previous years there was a mixture of panel review and mail-in-only review, all of last year's proposals were assessed within the context of the panel review, in keeping with comments made by the 2008 ASCAC COV that procedural consistency in proposal assessment was desired. As standard best practices, proposal PIs and co-PIs are precluded from panel participation, and all panel members are asked to identify potential conflicts of interest. Based on input from both the 2008 COV and the 2009 OA reviewers, greater transparency has been introduced into the INCITE program through more explicit definitions of expectations for leadership computing; the call for proposals clearly states that applicants must present evidence that their proposed production simulations can make effective use of a significant fraction—in most cases 20% or more—of the full configuration of the HPC systems requested for allocation. Further, the questions asked of the proposal reviewers are also now posted on the INCITE website, so authors can see the emphasis on leadership science. ALCF and OLCF staff members engage in several conference calls before the review process begins to discuss the general expectations for computational readiness of applications. Potential project authors are encouraged to request time on the system to carry out INCITE benchmarking or, for discretionary awards, to engage in porting and scaling activities in preparation for a future INCITE submittal. As described in Section 1 of this report, center staff members work with users and code developers year-round to support application development on leadership systems.

3.4.2 ASCR Leadership Computing Challenge Program

Open to scientists from the research community in academia and industry, the ALCC program allocates up to 30% of the computational resources at NERSC and the leadership computing facilities at Argonne and Oak Ridge for special situations of interest to DOE, with an emphasis on high-risk, high-payoff simulations in areas directly related to the department's energy mission in areas such as advancing the clean energy agenda and understanding the Earth's climate, for national emergencies, or for broadening the community of researchers capable of using leadership computing resources. The call for proposals will be issued annually for single year proposals; however, proposals for single year allocations may be submitted at any time during the calendar year. Proposals submitted to the ALCC program will also be subject to peer review of scientific merit based on guidelines established in 10 CFR Part 605.

3.4.3 Director's Discretionary Program

The DD program provides a valuable mechanism for the investigation of rapidly changing technology or unanticipated scientific opportunities that frequently arise outside the standard (INCITE) annual proposal cycle. The goals of the DD program are threefold: development of strategic partnerships, leadership computing preparation, and application performance and data analytics.

Strategic partnerships are partnerships aligned with strategic and programmatic ORNL directions. These are entirely new areas or areas in need of nurturing. Example candidate projects are those associated with the ORNL Laboratory Directed Research and Development Ultrascale Computing 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).

The DD program must help to identify and develop new computational science areas expected to have significant leadership class computing needs in the near future as well as exploit existing computational science areas where a leadership computing result can lead to new insight, an important scientific breakthrough, or program development. Candidates for such leadership preparation projects include those from industry, the SciDAC program, end station development, and exploratory pilot projects.

The DD program must also enable porting and development exercises for infrastructure software such as frameworks, libraries, and application tools; and support research areas for next-generation OSs, performance tools, and debugging environments. Candidates for such application performance and data analytics projects include application performance benchmarking, analysis, modeling, and scaling studies; end-to-end workflow, visualization, and data analytics, basic computer science research; and system software and tool development.

In 2009 the OLCF initiated the Industrial Partnerships Program as part of a new strategic goal within the DD program to promote usage of high performance computing across U.S. industry. For more details about this program and recent success stories, see Section 5.9.

The duration of DD projects is typically shorter than INCITE projects for two reasons: DD projects are intended to solve a problem within a finite period of time (e.g., scalability development) or be a prelude to a formal INCITE submittal, which is the appropriate vehicle for long-term research projects. The actual DD project lifetime is specified upon award, where most allocations are for less than 1 year.

The Resource Utilization Council (RUC, see Section 2) makes the final decision on DD applications, using written input from subject matter experts. Once allocations are approved, DD users are held to basically the same standards and requirements as INCITE users.

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

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

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

[a]All figures are percentages.

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 Gov 6.4 Industry 47.1 Academic
2009	47	51	125	38	55.9 DOE 0.7 Gov 9.9 Industry 33.5 Academic
2010	77	85	160	85	46.0 DOE 2.3 Gov 12.2 Industry 39.5 Academic

4. FINANCIAL PERFORMANCE

CHARGE QUESTION 4: How well is the program executing to the cost baseline pre-established during the previous year's Budget Deep Dive? Explain major discrepancies.

The total cost for FY 2010 was $96,114K. Of this 21.3% was spent on effort, 43.7% on lease payments, 10% on center charges (utilities), 10.1% on computer system maintenance, and 14.9% on other costs. The OLCF carefully managed costs in fiscal year (FY) 2010 to execute the FY 2010 OLCF operational requirements and meet the targeted system availability and number of hours delivered.

The OLCF carefully managed costs in FY 2010 to execute the FY 2010 OLCF operational requirements. The proactive management of budgets was aided by additional funds received late in FY 2009 and deliberate management of FY 2009 costs to maximize the carryover into FY 2010 in order to provide the necessary $96 million for FY 2010 expenses. The FY 2011 budget for OLCF, based on the President's requested funding of $96 million, will allow the OLCF to meet its scheduled improvements and ongoing obligations.

For OMB reporting, the OLCF is considered a mixed-life-cycle investment, with development, modernization, and enhancement (DME) as well as operational elements. The DME portion of the OLCF budget includes project costs related to bringing in a new computer system and the system lease costs (if any) prior to acceptance. After acceptance, all costs related to the systems are included in the operational portion of the OLCF budget. Therefore, as new systems are acquired, there is a continuous cycle of DME and operations within the OLCF program budget. The OLCF tracks all costs against the yearly budget in functional categories (leases, utilities, etc.) and cost types (labor, subcontracts, etc.) and by DME and operations. This allows the OLCF to monitor costs against planned budgets in numerous important ways. The OLCF is aided in this ability by a powerful SAP financial system that can pull information from the time-reporting system and the procurement system. The financial status of the OLCF is monitored daily by the OLCF finance officer and at least monthly by OLCF management.

4.1 FY 2010 ANNUAL FINANCIAL ESTIMATE, RISKS, AND BUDGET HIGHLIGHTS

The planned OLCF budget for FY 2010 was $95 million. The President's FY 2010 budget included $88 million for the OLCF. The OLCF actually received funds of $81.2 million.[*] After discussions with the ASCR program manager regarding this shortfall and ASCR's request to extend the operation of the Cray XT4 system beyond its planned shutdown in March 2010, it was agreed that funding of $2.3 million would be provided to the OLCF through R&E funds. These research funds were provided to cover work on performance tools and a limited amount of application readiness work. The significant funding shortfall of $14 million required the OLCF to postpone much of the OLCF-3 project, which had been planned to start in FY 2010, and to deplete the carryover from FY 2009. The value of the DME portion of the FY 2010 budget was limited to $512,000 for project planning activities.

ARRA funding of $19.9 million for the Six-Core Upgrade Project, received at the end of FY 2009, was fully costed with the completion of the project in December 2009.

The OLCF costs for FY 2010 have tracked closely to the FY 2010 budget. The OLCF yearly costs are dominated by costs in two areas: effort and system lease. The lease costs are locked in once the lease terms are negotiated and therefore will track to the budget as planned for each year. As planned, in 2010 $42 million will be costed against the Cray XT5 leases.

[*]Funding received through July 2010.

In FY 2010, the OLCF lost some staff due to high demand for certain application development and support disciplines. Additionally, two members of the OLCF management team charged more time to other projects than originally planned (Rogers—NOAA, Kothe—Consortium for Advanced Simulation of Light Water Reactors) and, in order to specifically conserve budget in a tight funding year, some planned new hires were not hired. This caused total expenditures for effort to be less than planned. Total OLCF full-time equivalent (FTE) employees for FY 2010 are 71.3. The planned FTEs for FY 2011 are 78; this includes a number of new or replacement hires in SciComp.

Because of the extremely tight budget in FY 2010, the OLCF used a formal "budget change process" for reallocating budgets from effort categories to other expenditures to maximize the effectiveness of these funds. A listing of opportunities was maintained and specific approvals were obtained for incremental procurements as the year progressed.

Financial figures are based on actual costs incurred as of June 30, 2010, and planned costs for July, August, and September. Table 4.1 summarizes project costs for 2010.

Table 4.1. FY 2010 Budget ($K)

FY 2010 Budget and Costs	**FWP**			**Total**
OLCF Program Budget:				
FY 2009 Carry Forward[a]			12,969	
New Budget Authority FY 2010				
Main OLCF Operations[b]	ERKJZN1	81,168		
Performance Tools	ERKJU41	1,544		
Application Development[c]	ERKJM02	2,000		
			84,712	
Total				97,681
OLCF DME Costs:				
DME Planning			512	
Total				512
OLCF Operating Costs:				
Effort			20,460	
Major Computational System Leases			42,000	
Center Charges			9,600	
Maintenance			9,703	
Subcontracts			3,805	
Center Balance Activities			3,426	
Other Hardware Purchases			4,214	
Other (travel, outreach, etc.)			850	
Performance Tools			1,544	
Total				95,602
Total OLCF Program Costs:				96,114
Carryover Planned[d]				1,567

[a]Uncosted/uncommited funds include FY 2009 end of year (EOY) funding of $2.880M.

[b]FY 2010 funding is based on Budget Authority received through July 2010 and is subject to change if additional funds are received.

[c]Expected funding of $2M not received as of July report date but included in table totals.

[d]With expected $2M EOY funding, carryover planned would be $1,567K.

A comparison between the planned cost categories and the actual costs for FY 2010 is reflected in Figure 4.1. Also included is the planned budget by category for FY 2011.

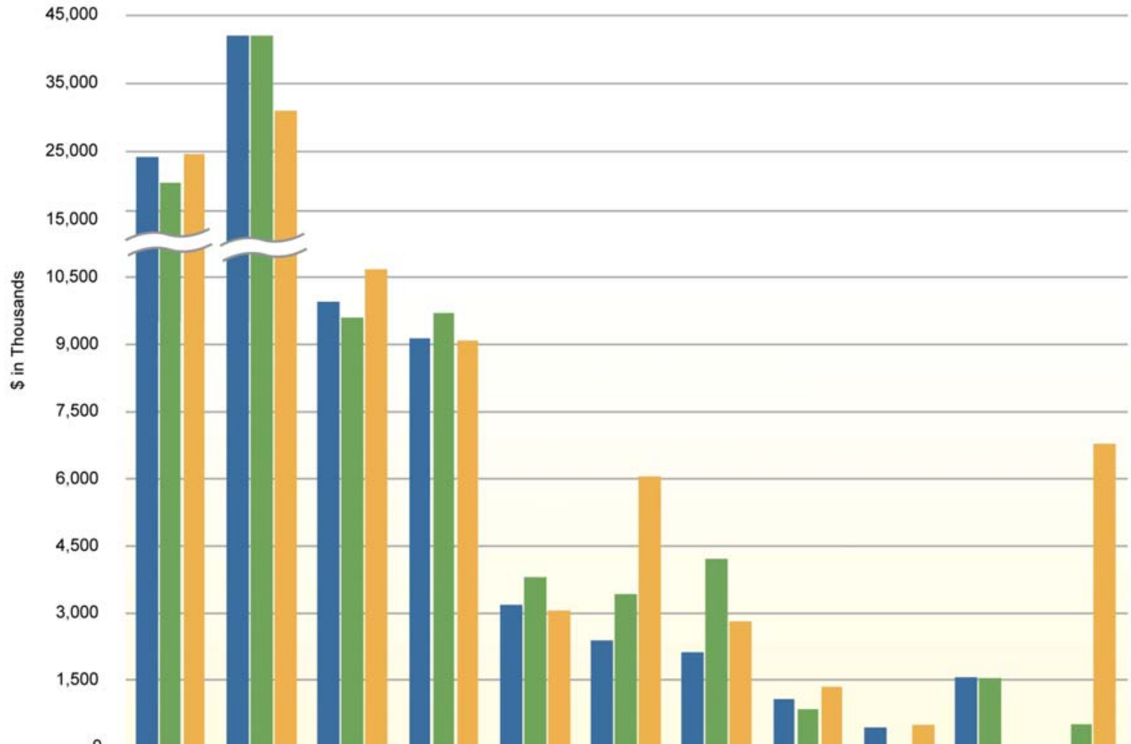

	Effort	Cray Leases	Center Charges	Maintenance	Sub Contracts	Center Balance Activities	Other Hdwre Purchase (HPSS)	Other (travel-outreach)	Reserve	AP Tools	DME Activities
FY10 Budget	24,241	42,000	9,955	9,135	3,187	2,391	2,125	1,075	441	1,564	
FY10 Cost/Commitments	20,460	42,000	9,600	9,703	3,805	3,426	4,214	850	0	1,544	512
FY11 Budget	24,658	31,000	10,681	9,086	3,059	6,060	2,817	1,351	500		6,788

Figure 4.1. FY 2010 budget vs costs and FY 2011 budget.

4.2 FY 2011 ANNUAL FINANCIAL PLAN

Financial figures are based on plans developed in July 2010. All projected costs include escalated rates based on the ORNL ePlan tool.

Table 4.2 summarizes the planned budget and costs for FY 2011.

4.3 FY 2011 BUDGET HIGHLIGHTS

The planned OLCF budget for FY 2011 is $96 million. Because the lease payment on the Cray XT5 drops from $42 million in FY 2010 to $31 million in FY 2011, there is a small amount of headroom in the FY 2011 budget to begin the OLCF-3 project previously planned for FY 2010. The DME includes plans for performance tools and advanced application development, required test beds, and minimum requirements for site preparation. Also, as expressed in discussions with ASCR, an alternative source of funding will be sought to continue operations of the Cray XT4 system for FY 2011 (currently included in the FY 2011 budget).

Table 4.2. FY 2011 Budget ($K)

FY 2011 Planned Budget and Costs	FWP			Total
OLCF Program Budget:				
FY 2010 Carry Forward[a]	Various		1,567	
New Budget Authority FY 2010				
Main OLCF Operations	ERKJZN1	96,000		
			96,000	
Total				97,567
OLCF DME Costs:				
Site Preparation			2,000	
Performance Tools and Testbeds			3,000	
DME Effort			1,788	
Total				6,788
OLCF Operating Costs:				
Effort			24,658	
Major Computational System Leases			31,000	
Center Charges			10,681	
Maintenance			9,086	
Subcontracts			3,059	
Center Balance Activities			6,060	
Other Hardware Purchases			2,817	
Other (travel, outreach, etc.)			1,351	
Management Reserve			500	
Total				89,212
Total OLCF Program Costs:				96,000
Carryover Planned				1,567

[a]Carryover includes late funding FY 2010 of $2M in FWP ERKJM02, but other funds which may be received after July 2010 are not included.

4.4 FY 2011 SPENDING RISKS

The new facility to house the OLCF-3 system brings uncertainty to the FY 2011 funding profile in the area of site improvements. We have very minimally budgeted for building related costs, but our estimate for site preparation varies depending on the options still under consideration for the computer center facility. Once FY 2011 funding is known and the facility option is approved, a more complete cost estimate and the funding requirement by year will be determined. Current estimates range from $2 million to $15 million.

Other funding uncertainty has added additional risk to the OLCF. It is anticipated that the OLCF funding will be part of a continuing resolution (CR) for at least 6 months. The budget has been planned to allow OLCF to push several budgetary items to the last 6 months of the year to provide adequate cash flow in the early months of the year. Figure 4.2 reflects the spending planned under a 6-month CR scenario.

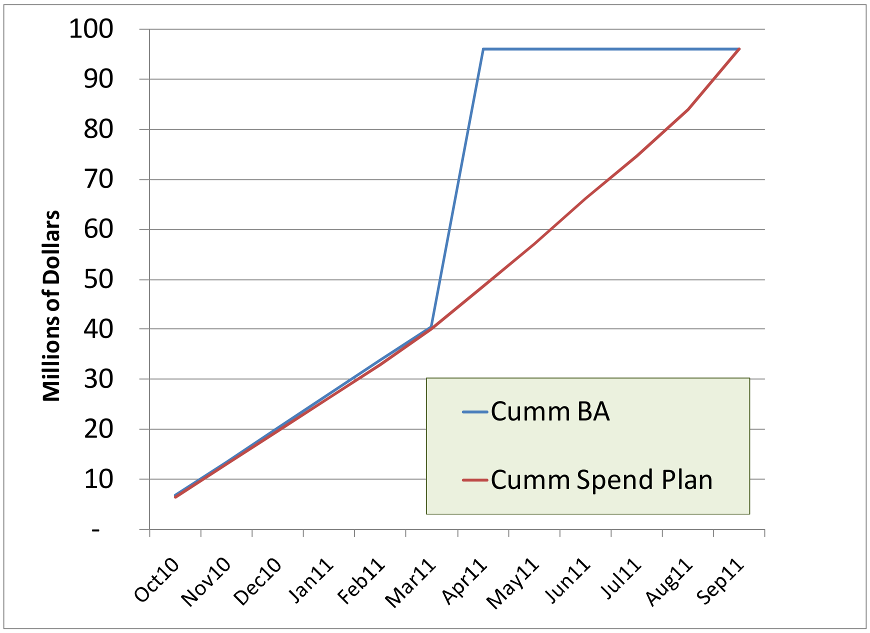

Figure 4.2. FY 2011 spending plan under a 6-month CR scenario.

If the planned funding of $96 million is significantly “shorted” or if the OLCF is required to operate at the FY 2010 funding level beyond 6 months, major budget cuts will be required. The reduced funding would potentially impact both operations and the OLCF-3 project activities. Should this occur, OLCF will work with ASCR to determine the correct strategy to overcome the shortfall.

5. INNOVATION

CHARGE QUESTION 5: What innovations have been implemented that have improved the OLCF's operations?

The breakthrough science taking place at the OLCF requires substantial behind-the-scenes support. While the supercomputers grab the headlines, their work depends on a robust infrastructure and scalable tools that include petascale file and storage systems, high-performance networks, scalable debugging and performance tools, and an array of other tools and enhancements. This section describes a series of high-impact innovations the OLCF developed and implemented this past year to support the leading edge scientific breakthroughs being pioneered on OLCF platforms and to advance HPC.

5.1 CENTER-WIDE FILE SYSTEM

A major innovation at the OLCF in 2010 has been the production release of Spider, the world's largest Lustre-based, center-wide file system.[*] A number of innovations developed in the process of productizing the Spider file system have benefited the OLCF and other centers. Performance improvements to the Spider parallel file system have enabled optimal performance on the underlying disk storage system. One such improvement was recently presented at the 2010 USENIX Conference on File and Storage Technologies.[†] Through the use of a novel file system enhancement known as asynchronous journaling, the performance of the DDN storage infrastructure for Spider was nearly doubled from 3 GB/s to more than 5.3 GB/s. To achieve the same level of performance in the absence of this software innovation, the DDN storage infrastructure would need to be expanded by more than 73% at a cost of more than $9 million. The performance benefits of this innovation have directly benefitted applications at the OLCF. For example, runtime was reduced more than 50% for the Gyrokinetic Toroidal Code (GTC) fusion plasma simulation.

Designed to support the peak bandwidth requirements of the OLCF's demanding workloads, the Spider parallel file system is capable of delivering remarkable aggregate performance. Figure 5.1 illustrates the

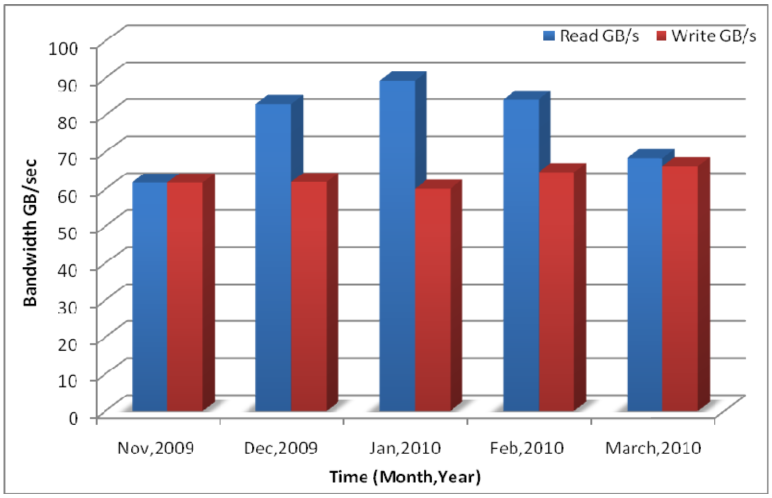

Figure 5.1. Maximum data rates on half of the storage controllers.

*G. M. Shipman et al., "The Spider center-wide file system; from concept to reality," Proceedings of the Cray Users Group, 2009; G. M. Shipman et al., "Lessons learned in deploying the world's largest scale Lustre file system," Proceedings of the Cray Users Group, 2010.

†S. Oral et al., "Efficient object storage journaling in a distributed parallel file system," FAST, 2010, pp. 143–154.

aggregate performance of half of the total storage controllers of the Spider parallel file system environment. These mixed workloads achieve nearly 75% and 60% of absolute peak read and write performance, respectively. This level of performance is seen under a mixed workload on multiple OLCF computational resources. Future improvements aimed at reducing network congestion will allow even higher levels of aggregate performance.

OLCF IMPACT ON INDUSTRY

ORNL's development efforts on the Spider file system have had a significant impact on Cray's I/O plans and have greatly influenced the scope of the I/O solutions Cray now offers its customers, ultimately resulting in a new product line for Cray. All parties involved with the Spider project—DDN, Sun, and Cray—have benefited from partnership with the OLCF. Based on OLCF development work, Cray included support of OpenFabrics InfiniBand and the Lustre router to Cray's service nodes for Spider. Cray now offers external Lustre as an I/O solution for customers who want the flexibility of sharing Lustre with other computational resources at their sites. For example, the Danish Meteorological Institute has an external Lustre file system shared between two Cray supercomputers in a configuration for high-availability computation. Cray has a number of systems planned for delivery this year that include external Lustre solutions, including the Hopper system at NERSC and a system at the Korea Meteorological Administration.

While ORNL has code-named the external Lustre file system Spider, Cray has opted to productize the Spider file system as esFS. It has the unique ability to serve multiple computing platforms simultaneously with high performance and availability. Before esFS, each Cray XT computing platform had its own file system, creating "islands" of data. Users of large-scale Cray XT platforms frequently undertake post-simulation data analysis and visualization on separate compute platforms. These workflows required costly data transfers from the simulation platform file system to the data analysis and visualization platform file system. The esFS eliminates these costly data transfers and redundant file systems by providing accessibility to multiple compute platforms at extremely high performance. Before esFS the parallel file system environment on the Cray XT platform was accessible only when the Cray XT platform was available. By decoupling the file system from the Cray XT platform, esFS provides file system accessibility independently of the Cray XT platform. In addition to eliminating islands of data, esFS allows storage systems to grow independently of the computational environment and thereby take advantage of differing technology curves for storage and compute technologies. The esFS solution can service a wide variety of demands, from modest installations with a few terabytes of capacity to world-class computing centers requiring petabytes of storage and hundreds of gigabytes/second of file system throughput.

A number of advances in large-scale network and parallel file system architectures were required in the development of esFS, including coupling of differing advanced network technologies, engineering of system reliability at all levels, and improvement to file system scalability to support tens of thousands of file system clients. ORNL's work on decoupling the Lustre servers from the Cray supercomputer has made Cray's I/O solutions more flexible and more easily accessible.

Within the next year, Cray expects a number of third-party vendors to provide Lustre appliances. "The development of this external Lustre business is a result of ORNL's technology leadership and the efforts of the engineers at ORNL, Cray, and Oracle who made this project a success, providing a fully functional, shared, scratch file system that supports the world's most demanding high-performance computing environment," says John Carrier, senior software engineer at Cray.

Moving toward a centralized file system requires increased redundancy and fault tolerance. Spider is designed to eliminate single points of failure and thereby maximize availability. By using failover pairs, multiple networking paths, and the resiliency features of the Lustre file system, Spider provides a reliable high-performance centralized storage solution greatly enhancing our capability to deliver scientific insight. One of the many fault-tolerance techniques used is dm-multipath, available within Linux. This allows Spider to sustain I/O and provide file system resources to the compute platforms during a failure of one or more components. DDN controller failures, InfiniBand cable failures, or InfiniBand host channel adapter failures are all examples of modes where performance will be degraded, but the file system will continue to function. In other implementations without dm-multipath, any one of these failures will cause an outage for all compute platforms that use the file system. This also allows the upgrade of DDN firmware without an outage. These innovations have improved the availability of the Spider file system by allowing the file system to remain accessible in the event of these failures, allowing us to achieve 99.7% scheduled availability.

In addition to developing and deploying next-generation file system technologies, the OLCF has developed core competencies in the operational management of these systems. In late February 2010, multiple storage hardware failures resulted in corrupt data on two RAID (redundant arrays of inexpensive disks) sets. Working around the clock, OLCF staff in HPC Operations and TechInt identified files potentially impacted by the event (less than 0.3%), temporarily restricted access to these files, and brought the file system back online to allow applications to continue to run on the OLCF computational platforms. Recognizing the impact of data corruption on our users, OLCF staff developed a set of tools to identify and isolate individual files with corrupted data using block-level analysis. Using these tools, the OLCF narrowed the number of files with known data corruption to 0.03% and files unlikely to be impacted but requiring more analysis to 0.11%. These tools allowed individual files impacted by this event to be identified and UAO to contact owners of these files with information about the impacted files. Without the core competencies in parallel file system technologies developed at the OLCF and our coordinated response to this event, an eventual outcome would likely have been the formatting of the file system and the loss of more than 285 million files. The lessons learned from this event have since been shared with other facilities, including ALCF, LLNL, and NERSC. Tools developed in response to this event are being further developed and will be released as open source software to benefit the larger parallel file system community.

5.2 APPLICATION TRACING, ANALYSIS, AND OPTIMIZATION

Among the challenges of deploying leadership computer systems is encountering problems no one has seen before. This means that no one has developed the tools to deal with these problems. As a result, the OLCF has developed an integrated system-monitoring framework, with customized tools, that collects system events and aggregates performance data. This framework allows for detailed analyses of system failures and aberrant performance as well as root cause analysis of system interrupts.

5.2.1 Tracing and Trace Analysis

One of the greatest challenges we have encountered is tracking down aberrant application behavior that may be completely benign at smaller scales but pathological at leadership scales. Quickly tracking down the source of these issues at scale using traditional approaches is intractable. To address this, we ported a number of tools to the OLCF environment, while others were developed and are now maintained by the OLCF.

One such tool developed by the OLCF was MDSTrace. MDSTrace allows the OLCF to monitor load on the Spider file system and correlate load with specific applications. While there are a number of tracing tools in the HPC community, no existing tool provided this functionality. MDSTrace works by capturing a 60-second “slide” of remote procedure call traffic to the Lustre metadata server every 10 minutes and then generating a report detailing what the server was doing, which applications were running, queue times, and how long it took to perform each action. This tool allows staff to identify applications that create substantial load on our parallel I/O environment in an automated fashion.

The unique ability of the MDSTrace tool to identify applications that are causing or contributing to heavy MDS load has benefit to the broader HPC community. We are now working to further harden MDSTrace and plan on releasing this tool to the broader community as an open source software project.

In some cases, rather than develop completely new tools, we have been able to modify existing open source software for the HPC environment. A good example of this is strace, the popular Linux diagnostic and debugging tool. We have developed methodologies to automatically start strace within compute jobs to collect detailed information about application behavior. We have also developed utilities to parse the collected traces to detect several common pathological behaviors that generally result in performance degradation. We collaborated with colleagues at NCSU who developed ScalaTrace, an open source tool for tracing MPI (message passing interface) events, to customize it for the Cray XT environment. Our work in porting and optimizing these tools on Cray platforms benefits other sites that use Cray systems, such as LBNL and SNL, as well as the broader HPC community. These tools have been critical in identifying the root cause of application performance degradation.

Using the tracing tools discussed above, OLCF staff can pinpoint applications potentially causing performance problems and work with the project liaisons (Section 1.2) to improve application I/O to both improve runtimes and reduce load on the parallel I/O environment.

5.2.2 DDNTool

The Spider file system is built around storage controllers from DDN. The controllers have an application programming interface (API) that allows performance and fault information to be queried over the network; however, querying the 96 separate controllers in the file system individually is impractical. To address this, the OLCF has developed a utility called DDNTool to regularly query the controllers and aggregate the results. This information is stored in a MySQL database. Because client programs all connect to the MySQL server, multiple users can monitor the system without the risk of overloading the DDN hardware and slowing performance.

In the last year, DDNTool has been used to query system temperature sensors (14 sensors per DDN controller; 1,344 total) and display information for those whose temperature was above a specified value. It has also been used to capture historical performance data such as storage system bandwidth usage and aggregate I/O operations per second and I/O request size distributions. Currently DDNTool is being used to capture performance and fault data for display on web pages showing read/write peaks on a per-minute basis and to generate graphs showing the delay in committing SCSI transactions to disks.

DDNTool has allowed us to build new systems for monitoring our storage system infrastructure. Figure 5.2 illustrates one such use of DDNTool: monitoring a variety of file system metrics and performance data.

Because of common requirements at other facilities that use DDN storage controllers, we are now working to release DDNTool as an open-source software project.

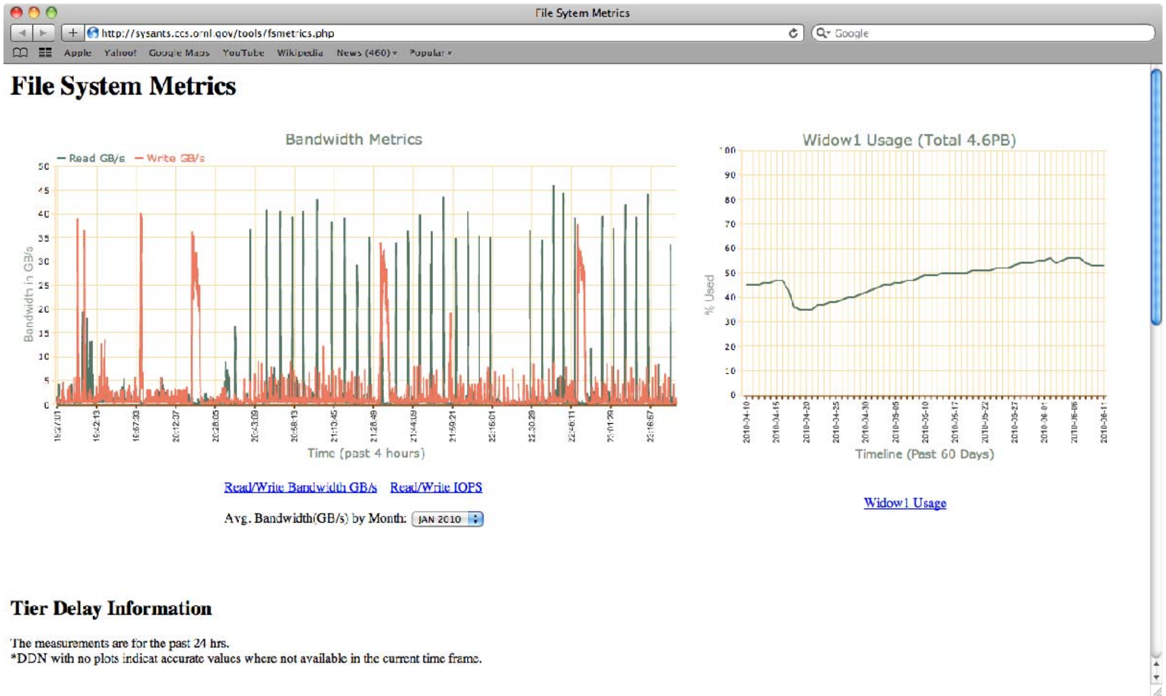

Figure 5.2. File system metrics.

5.2.3 System Log Analysis

Log data analysis is one of the system administrator's main tools for understanding system failures and isolating problems; however, deriving meaningful, actionable information from the deluge of log data generated by modern leadership computing platforms is a complex and tedious task. The sheer volume of log data can be daunting—for example, a single event on Jaguar can generate a few hundred thousand log entries in less than a minute. Much of this information is redundant, and log messages are often highly unstructured, further complicating the task of log data analysis.

To address this problem, we have devised a three-step process to present log data in a more concise format, without loss of information, and thus render it more amenable to analysis.

1. Log messages are preprocessed using knowledge of the structure of the error messages gained from the source code.
2. Log messages are clustered within a time window using a common feature set. This provides a precise summary of events for a given timeframe and an indication of the system state and environment.
3. Clustered log messages are then correlated with hardware and associated with job information to find trends tied to individual applications.

This technique reduces the total volume of log data while preserving and enhancing its diagnostic value, thus providing system administrators with meaningful information in a concise format for root cause analyses. Figure 5.3 illustrates a time series view of system log messages and the ability for our system log analysis utility to represent a series of correlated events as a single entity, thereby dramatically reducing the volume of data that the administrator must review.

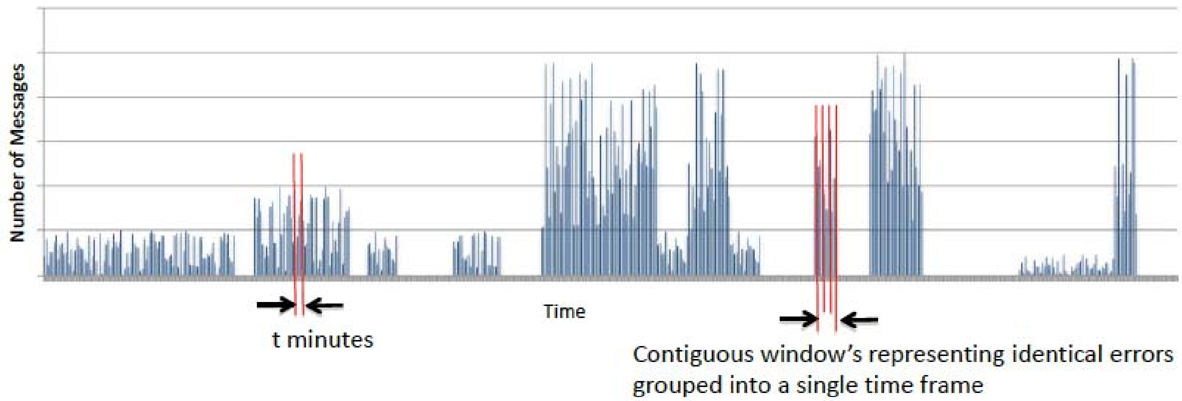

Figure 5.3. Time series view of console log messages.

This technique has been fully prototyped and tested against offline datasets; using it we have been able to demonstrate a high correlation between applications and specific error messages. Future development efforts will concentrate on processing log streams in real time to enable near-real-time response to system failures and the presentation of these failures in a meaningful and concise format to facilitate a shorter time to resolution.

5.3 PARALLEL DATA TOOLS

The Cray XT5 Jaguar coupled with the Spider parallel file system provides the ability to create multiple terabyte datasets, and indeed, some current OLCF users have datasets that exceed 10 TB. Given that the amount of data generated on our systems is growing exponentially, it is not unreasonable to expect that we will see 100 TB datasets in the not-too-distant future. However, as we continue to scale our computational platforms, we are finding that the use of traditional serial tools for managing data, such as tar and cp, presents a challenge to productivity. For example, cp and tar move data at an average rate of 265 MB/s. For a 100 TB target dataset, the time for such serial tools to move data would be about 9 days. OLCF staff believe this is unreasonable—particularly when the datasets can be created in as little as a day on Jaguar. Efficient management of these datasets requires dramatically improved system utilities that preserve striping metadata and are parallelized to exploit the benefits of large scale parallel file systems.

Consequently, we have developed a parallel copy utility, spdcp, and a parallel archive utility, pltar.
In standard benchmarking tests, spdcp has been shown to provide dataset copies up to 93 times faster than the standard Linux cp utility. The parallel archive utility (pltar) has demonstrated up to 49 times the speed of the GNU tar utility.

Work on spdcp is complete for the most part, and it has been released under general public license for the benefit of the larger HPC community and added to the Lustre User Toolkit. pltar is still at the prototype stage; however, we will complete our production release this year and make this available to the broader community via an open source license.

5.4 A FRAMEWORK FOR INTEGRATED END-TO-END SCIENTIFIC DATA MANAGEMENT TECHNOLOGIES FOR APPLICATIONS

To support the OLCF goal of accelerating scientific discovery, we not only have to maintain the OLCF HPC platforms, we also must develop tools that allow users to generate fast, scalable, metadata-rich I/O and perform the mundane tasks needed to produce meaningful scientific results. The OLCF approach to this challenge is called FIESTA, the Framework for Integrated End-to-End Scientific Data Management

Technologies for Applications. The two main elements of FIESTA are an adaptable I/O system (ADIOS) and the eSimMon dashboard (accompanied by the Kepler workflow system).

5.4.1 ADIOS—ADapatable I/O System

Currently the ADIOS API (developed jointly by the OLCF and colleagues at the Georgia Institute of Technology, SNL, Rutgers University, NCSU, and Auburn University) is used daily by many of the major scientific codes running on HPC platforms, including XGC1, XGC0, GTC, GTS, GEM, M3D, and S3D. ADIOS provides a simple, flexible way for scientists to describe the data in their codes that may need to be written, read, or processed outside of the running simulation, thus increasing productivity. Used with the GTC code, ADIOS achieves more than 60 GB/s when writing data on the Cray XT5, and high read performance for analysis and visualization has also been achieved—more than 35 GB/s on the Cray XT4. ADIOS 1.0 was released as open source software in November 2009 and is the version currently in production mode. ADIOS 1.2 was released in July 2010. The most notable features of the new release include asynchronous buffered I/O with staging, more statistics in the ADIOS binary packed file format, and the creation of a new method for even faster I/O on the Lustre file systems.

5.4.2 Dashboard—Electronic Simulation Monitoring (eSimMon)

Running simulations is only half the story. Once the simulations have been run, the results must be analyzed and visualized. Petascale simulations that require massive amounts of computing power generate massive amounts of data—and as the complexity of the simulations increases, so does the complexity of the results. Add to this the highly collaborative nature of science—projects frequently include teams of specialists from multiple fields, all of whom want to view and analyze the data from their unique perspectives—and you have all the ingredients of a perfect storm. In the case of the OLCF, most users are remote, and transferring the massive amounts of data generated in simulations is becoming increasingly impractical. Thus there is an increasing need to provide scientists with the tools to analyze and visualize their data remotely and in real- or near-real-time, which enables early insights; allows corrections and/or adjustments on the fly, helping to avoid potential errors and thus increasing efficiency; facilitates identifying and zooming in on regions of interest; and because the power of high-performance computers is used in preprocessing, streamlines/simplifies data transfer and data transfer needs.

All of these factors drove OLCF development of the eSimMon dashboard, a web-based interface that allows scientists to access currently running simulations, launch new simulations, and revisit/analyze the results of existing and old runs, providing comprehensive end-to-end data and workflow management. The eSimMon dashboard gives users access to a wide variety of useful information, including how many jobs are currently running on particular platforms, performance text files, interactive images and movies generated from simulations, etc. (Figure 5.4). With eSimMon, scientists can even add notes to data files, giving collaborators the ability to share thoughts and ideas about the science as it is being created. Because the dashboard was developed with user ease as a requirement, a lot of the minutia of data manipulation is handled transparently, leaving scientists free to concentrate on results, regardless of individual computing abilities. While analytical/visualization packages have been incorporated in eSimMon, a plug-in capability allows users to access external analysis and visualization packages.

eSimMon has been tested and refined in complex fusion and combustion simulations over the past 2 years and has allowed researchers to monitor their simulations (with the help of the Kepler workflow automation tool) and analyze and visualize their data collaboratively. Currently eSimMon is being used on the same projects as ADIOS, and the two tools operate synergistically to increase productivity and user satisfaction. We are preparing to release eSimMon 1.0 as open source software in November 2010, adding an important element to the HPC toolkit available to scientists and engineers.

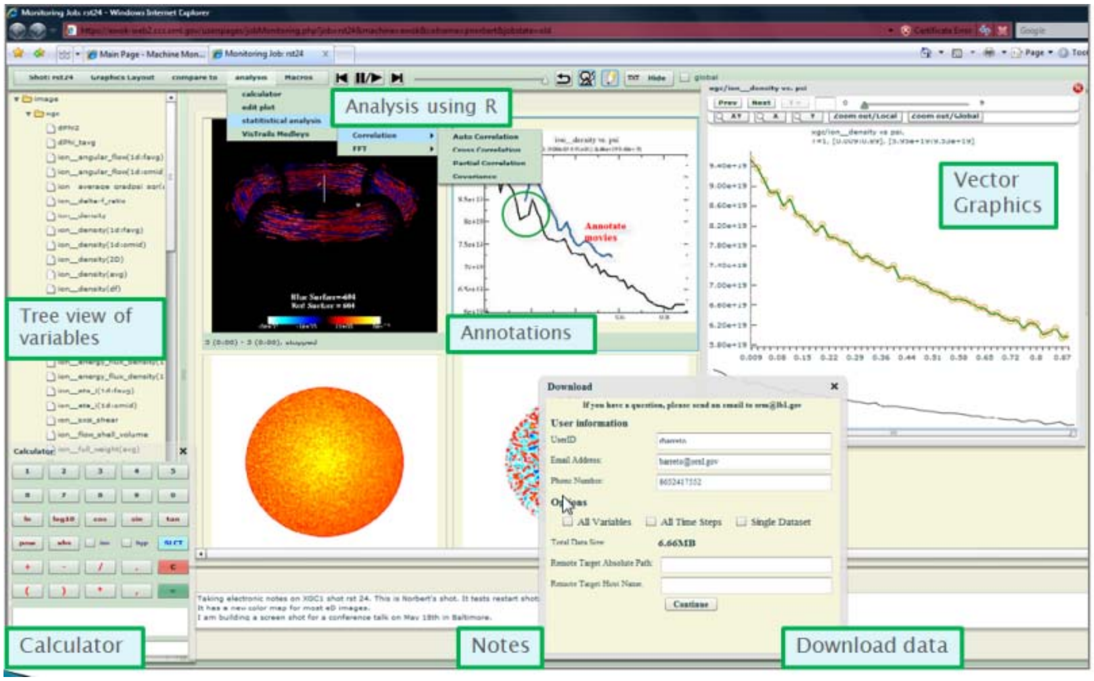

Figure 5.4. Screenshot of typical eSimMon view.

5.5 CENTRALIZED MAINTENANCE OF SOFTWARE

The OLCF maintains numerous software packages, libraries, tools, and applications to support research and improve the productivity of the computing environment. While most computing facilities find software management a challenge—name and version control are just two of the issues involved—at the OLCF, with hundreds of packages running on multiple machines, the problems can be staggering. For most of this software, the OLCF supports multiple versions and multiple builds. Because awards of time on OLCF computers are made to users across scientific disciplines, there is little homogeneity among the various supported applications.

Like most computing centers, the OLCF uses file management systems; however, most that are commercially available lack the degree of automation and control needed. User ease is another consideration—one the OLCF takes very seriously. In this last respect, the current trend is away from users having to know which directory-specific files and software are on systems and toward modularization—another way of automating the whole process and improving the user experience. Cray and other vendors already provide a lot of software in modules (modulefiles), so the concept is one with which users are already familiar.

To address these issues, the OLCF has developed SWTools, a centralized framework and process for building, deploying, and updating software packages.

One unique and useful feature of SWTools is that it automatically builds web pages for the software running on our systems, including version numbers, locations, and testing information, so that when software is updated, the associated web pages are automatically updated (Figures 5.5 and 5.6).

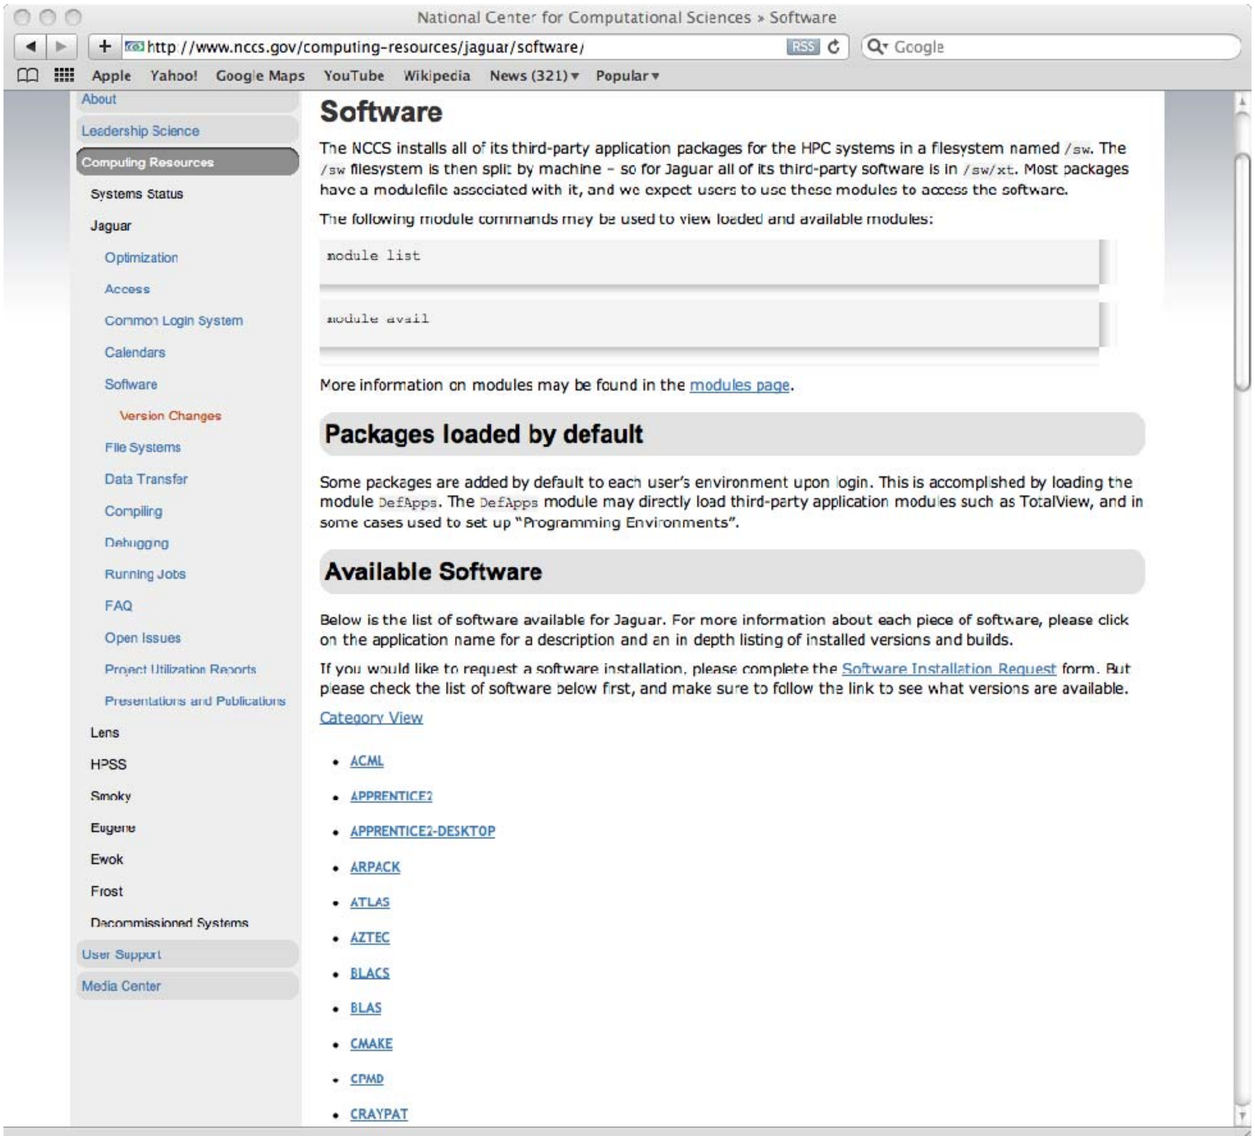

Figure 5.5. Software list web page automatically generated by SWTools. Each bulleted name is a link to a web page explaining how to use the software and what versions and builds are available. The available software can be viewed in alphabetical order, as shown, or in category view, in which software is categorized by its function (e.g., libraries, tools, science applications, etc.). When new software is installed, SWTools can automatically regenerate this web page to include the updated information.

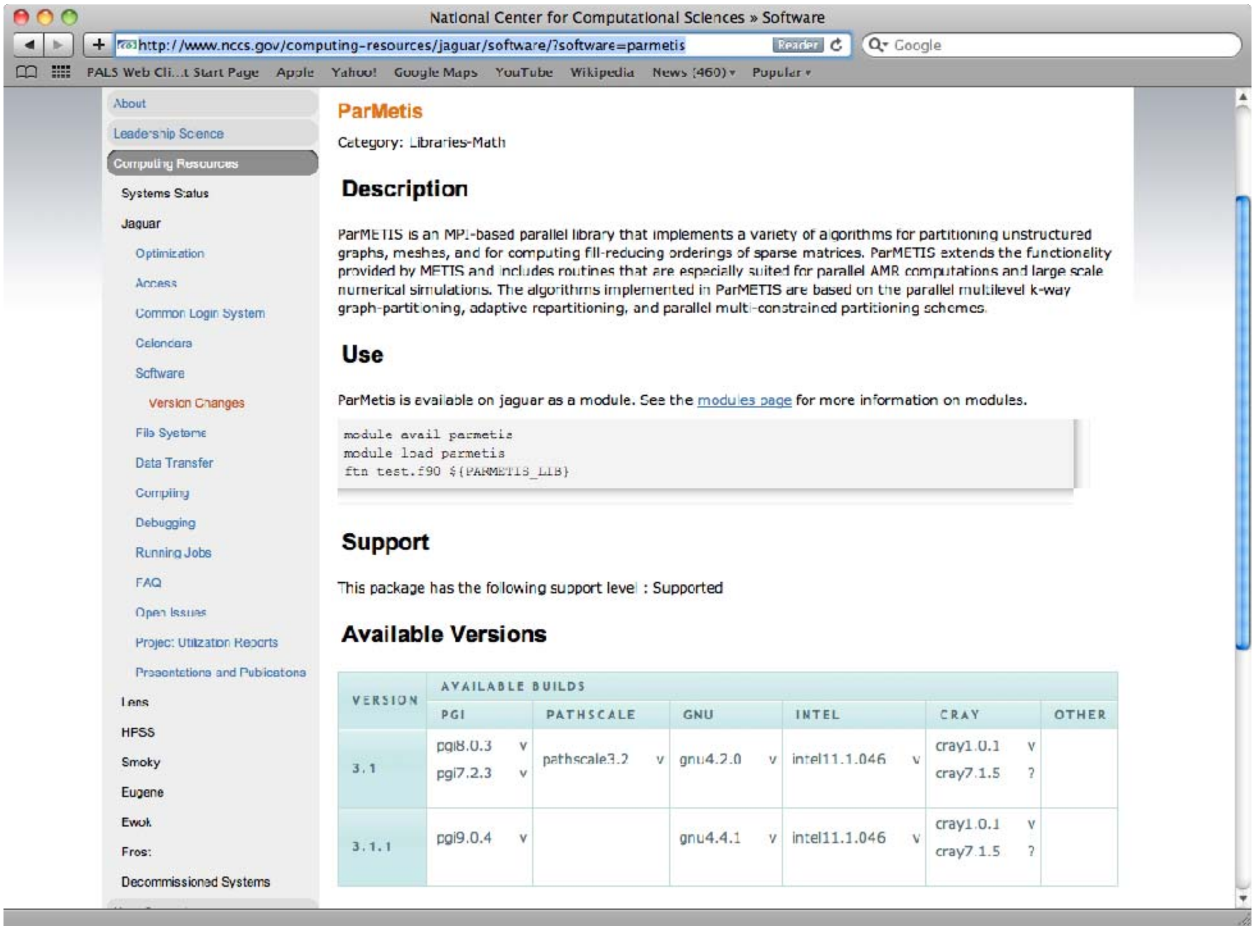

Figure 5.6. Software web page. For each available package, a web page is automatically generated, with a brief description, an explanation of how the software can be used, the support level of the package, and version and build availability information. When new versions and/or builds are installed, SWTools can regenerate the page with the updated information.

System administration is not limited to naming, loading, and testing software but also includes maintaining, updating, and determining when to deprecate and/or remove a particular software package. Without a way to methodically track usage, much of this aspect of system administration is subject to guesswork. To address this issue, the OLCF has developed a database and reporting system to transparently track all software packages installed on our systems. The system tracks how often a particular software package is used, who uses it, and even how it is used. Because it tracks software versions, we can check with users to determine why particular versions are being used and steer them to newer versions with possibly greater advantages, thus improving the user experience—and ultimately results. Another benefit is that when we detect problems with a piece of software running on our systems (see scientific liaison work described in Sections 1.2.1 and 1.2.2), this tool allows us to quickly and easily determine the identity of all its users so we can contact them and suggest alternatives or perform the necessary upgrades. Ultimately, use of this software enhances system administration decisions, leading to greater efficiency and cost savings as well as a better user environment.

Since implementing SWTools and the tracking system, we have been able to do the following:

- anticipate user needs and requests more efficiently;
- ensure a common, up-to-date software environment on all of our computing resources;

- provide a more customized, automated user environment;
- determine and change defaults more effectively and efficiently;
- provide user information more effectively and efficiently; and
- control and maintain all our software resources more efficiently and effectively.

We have discussed SWTools in workshops and similar venues, where its uniqueness and potential value to the whole HPC community have been recognized. Therefore, SWTools is being prepared for release as an open source package, which will benefit the broader HPC community. Currently we are working on user documentation and a user manual and hope to release SWTools in the next year.

5.6 HIGH-PERFORMANCE STORAGE SYSTEM

HPSS refers both to the software for managing petabytes of data on disk and robotic tape that was developed through an industry/national laboratory collaboration and to the long-term storage system using that software here at the OLCF. HPSS software provides highly flexible and scalable hierarchical storage management that keeps recently used data on disk and less recently used data on tape. It has become the de facto standard for HPCFs and any facility with massive storage needs.

This past year, the OLCF HPSS team has been involved in several high-level HPSS projects, described in the following sections.

5.6.1 HPSS Version 7.3

HPSS version 7.3 was released in 2010, and the OLCF is currently in the process of converting to it. The upgrade is being approached in a stepwise fashion so that the required upgrades to the system infrastructure can be made to ensure the integrity of files and SA. The release incorporates a number of major improvements for users, including the following.

- Better performance for handling small files. For most systems it is easier and more efficient to transfer and store big files; these modifications made improvements in this area for owners of smaller files.
- Tape aggregation. The system will now be able to aggregate hundreds of small files to save time when writing to tape—a huge gain for the OLCF.
- Multiple streams or queues of what HPSS refers to as “class-of-service changes.” This will enable the system to process multiple files concurrently and, hence, much faster, another huge time saver for the OLCF and users.
- Dynamic drive configuration. The configuration for tape and disk devices may now be modified without taking a system down, giving the OLCF tremendously increased flexibility in fielding new equipment, retiring old equipment, and responding to drive failures without affecting user access.
- User-defined attributes. In response to user requests, the system will now provide a way to add application-specific information such as checksums or expiration dates to files, a great enhancement for users.

5.6.2 HPSS Version 8.1

To meet the needs of next-generation archival deployments, the HPSS collaboration is currently developing HPSS version 8.1. In this effort, HPSS will undergo a massive “rearchitecturing” to improve scalability, performance, and manageability for the coming exascale environment. Bigger machines, generating more data, will need exponentially larger storage systems. Among the challenges confronted by the development team are infrastructure and space needs, backup times for huge files, database structure for handling such massive amounts of data, and file integrity. Meeting these challenges will require partitioning of the database and the total rearchitecture of the HPSS core server so it can work with the partitioned database. As a core contributor to HPSS, ORNL participated in development of the exascale development document. We are working within the collaboration to develop 8.1 and within the larger HPC community to drive planning through 2016 and longer-term requirements analysis through 2020.

Since joining the collaboration in 1993, ORNL has been responsible for the storage system management (SSM) portion of the product and is the primary developer of the HPSS SSM interface. In addition, the OLCF has taken over primary development of a number of important HPSS subsystems during the past year, including the bitfile server (one-third of the core server) and the logging and accounting subsystems. To meet user needs and improve the user experience here, the OLCF has provided a leading role in the development of the following innovations, which are becoming standard in the HPSS community: real-time monitoring, dynamic device configuration, and project accounting.

5.7 IMPROVING APPLICATION PERFORMANCE THROUGH OPERATING SYSTEM SCALABILITY

In the past year, a high degree of variability in the performance of some applications was observed on the Cray XT5 system. While working to determine the source, TechInt staff discovered higher levels of OS activity on the Cray XT5 than on other HPC platforms. These high levels frequently interrupted the CPU, hindering maximum effectiveness of the affected applications.

It is well known that OS noise, defined as “interference generated by the OS that prevents the CPU from performing “useful work,” can cause desynchronization (jitter) in collective communication tasks. Recent research has shown that OS noise can cause serious variations in and overall degradation of parallel application performance—much like that observed on the Cray XT5 here and on other Cray XT series platforms at other HPC facilities. Further, such behavior is more pronounced in large-scale applications using certain types of tightly coupled communication primitives.

The OLCF initiated a collaborative effort with Cray to address the problem at leadership scale. The problem was isolated to numerous sources of noise in the Cray Linux Environment (CLE) architecture and infrastructure (relics of the OS development/customization process such as generic timer events and kernel daemons). The Cray-OLCF team weighed various options to address the problem and determined that a reduced noise kernel would be the most practical, providing a common Linux environment coupled with the scalability requirements of leadership computing. The prototype kernel developed by Cray and tested at ORNL aggregates most OS noise sources onto a specific core on each compute node and provides a user-controllable mechanism to prevent scientific applications from running on this core. Tests to date on applications with tightly coupled communication have demonstrated performance improvements greater than 30% at scale (Figure 5.7), and performance variability of collective communications at leadership scale has been significantly reduced.

While this solution has been tried on similar OS noise problems in the past, it has not to our knowledge been tried at this scale on Cray XT platforms. Cray plans to release this solution with future systems; therefore, as a result of OLCF leadership and this collaboration, CLE scalability has been further improved for all Cray XT series systems. According to Barry Bolding, vice president of the Scalable Systems Business Unit at Cray, “The development of Core Specialization benefitted greatly from the early collaboration with Oak Ridge National Laboratory. As soon as Cray’s Software Development

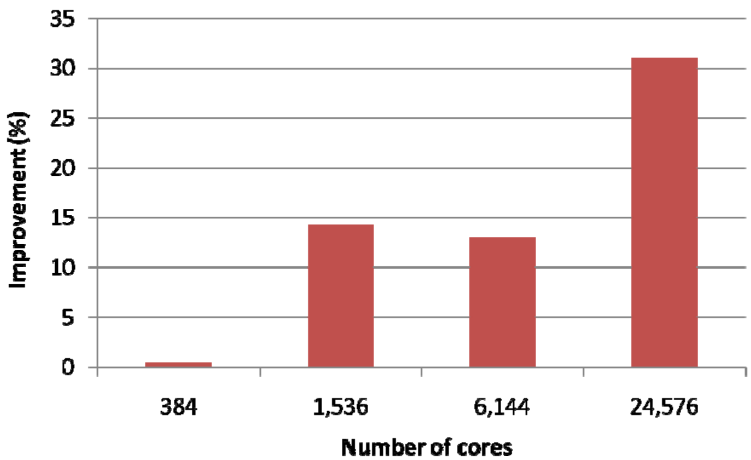

Figure 5.7. Completion times and efficiencies achieved on the Parallel Ocean Program running with the reduced noise kernel (compared to the stock kernel).

organization had a prototype available, ORNL installed it on Jaguar and carried out thorough performance evaluations. The ability to assess the impact of Core Specialization for real applications at very large core counts was invaluable for Cray, and would not have been possible without the collaboration with ORNL. The data from these tests allowed us to materially improve the first release of Core Specialization."

5.8 SCALABLE DEBUGGING AND PERFORMANCE TOOLS

It's no exaggeration to say that everything is bigger (and more complex) at the petascale. The task of debugging code certainly follows this rule. Debugging code running on massively parallel systems requires expert programming skills—skills that many OLCF scientific users don't have and are not expected to have. The OLCF model for HPC is to enable and maximize the abilities of scientists using our systems to produce excellent science.

However, the OLCF Application Performance Tools Group has long recognized that even with expert programming/debugging skills, we are reaching the limits of the current generation of debugging tools. Recognizing that it needs to provide users with practical tools to debug their programs, the OLCF initiated the Debugger Software Enhancement program in 2009 to scale current debugging tools to petascale production grade. Based on user needs, an understanding of the current state of HPC, and projections of where the technology is going, the OLCF developed very precise specifications and a phased approach for developing the next generation of debugging tools. In a competitive bidding process, Allinea Inc., the U.S. subsidiary of Allinea Software Ltd., was selected to partner with the OLCF on the 3-year project.

Divided into three phases, the project focuses on scaling Allinea's DDT, a comprehensive graphical debugger known for its ease of use. Work has begun on scaling the DDT debugger to run on 132,000 processors. So far, the debugger has demonstrated

- the ability to start up on 220,000 processes with start-up times similar to that of MPI jobs,
- scalable trace back and variable data collection (at 131,072 processes, stack traces and single variables are collected in less than 1 second),
- scalable attachment to running processes,

- a nonblocking and responsive graphical user interface that does not block while waiting on data to arrive and setting or clearing breakpoints, and
- scalable process control at 131,072 processes.

The project is ahead of schedule and already demonstrating unsurpassed scalability; current benchmarks show that key functions can be performed in a few hundred milliseconds or less—invoking each of the 100,000+ processes as well as reporting data back to the GUI front end—a vast improvement over anything previously possible at this scale. The first phase of the work is expected to be completed by summer 2010 and released in Allinea's product version of DDT, so impacts will be felt throughout the HPC environment. Graham believes, based on phase one results, that the debugging tool developed by project's end will be an important productivity tool.

In a separate-but-related aspect of this project, UAO staff members are working with Allinea on a pilot program to provide onsite training for using DDT. Initially UAO staff members will be trained, but plans include extending training to the broader OLCF user audience, with UAO staff members providing the training. The advantages to this approach are twofold: (1) we'll be able to provide application-specific training by trainers familiar with both our systems and user projects and (2) it will be more cost effective than relying on Allinea trainers. [Note: Positive results are already being seen in terms of user support (Section 1).]

As with debugging tools, the tools for performance measurement and analysis in the HPC environment either do not exist or are extremely limited in their capabilities. Consequently, users improvise solutions to meet their needs or, worse, ignore performance entirely—highly problematic in terms of efficient, effective use of the resources. The OLCF is addressing this issue head-on through collaboration with the Technical University of Dresden (TUD) to develop advanced performance analysis tools for the HPC environment. TUD was selected as a research partner for a number of reasons including their strong research arm and their Vampir suite of performance tools, which were judged to be appropriate for OLCF needs. We have purchased a Vampir license and loaded the suite, consisting of VampirTrace (including Beta CUDA support on Lens), Vampir, and VampirServer, to our platforms (Jaguar, Lens, Smoky). We have also trained UAO and SciComp staff in its use (see Section 1 for positive customer support results). Traces have been generated for application runs with up to 30,000 processes, resulting in trace data of up to 950 GB. Traces of that size were analyzed using up to 10,000 VampirServer processes and connecting into this VampirServer instance from the scientist's desktop computer. This tool suite is being targeted for use on the OLCF-3 platform. Understanding the performance of applications at scale is critical to putting together a credible plan to improve application performance at such scales. This work aims to address the difficult problem of providing detailed performance analysis for extreme-scale application runs.

5.9 CREATING A NEW INDUSTRIAL REVOLUTION THROUGH INNOVATIVE PARTNERSHIPS

In today's complex and highly interconnected global economy, modeling and simulation with HPC has become a critical ingredient in the recipe for industrial success. The OLCF is helping to create a new industrial revolution with its innovative HPC Industrial Partnerships Program, which enables companies to supercharge their competitiveness through access to world-class computational resources and domain expertise not available elsewhere (Table 5.1).

The program is a triple win, benefiting ORNL, industry, and the country. ORNL benefits by linking to some of the best thinking in corporate America, where many firms are pursuing scientific challenges similar or complementary to those at the lab in their quest to develop innovative products and services.

Table 5.1. Current Industry Projects

Corporate Partner	Program	Description
Boeing	INCITE	Development and correlation of computational tools for transport airplanes
General Motors	INCITE	Electronic, lattice, and mechanical properties of novel nanostructured bulk materials
Ramgen	ALCC	High-resolution design-cycle computational fluid dynamics analysis supporting CO_2 compression technology development
BMI	DD	Class 8 long-haul truck optimization for greater fuel efficiency
United Technologies Research Center	DD	Nanostructured catalysts for water-gas shift and biomass reforming hydrogen production
United Technologies Research Center	DD	Surface tension predictions for firefighting foams
GE Global Research	DD	Unsteady performance predictions for low pressure turbines

This cross-pollination of ideas often helps shed new light on laboratory research. Companies benefit through access to world-class HPC resources not available to them internally (or elsewhere) and also to our domain experts (Section 1.2) as they tackle complex, strategic problems that, if successfully addressed, will help them leapfrog their competition. As ORNL and companies advance together in their problem-solving abilities, they strengthen the nation's capacity to address national "grand challenges" like security threats, climate change, and energy consumption.

USING HPC TO ENGINEER SMARTER TRUCKS

BMI Corporation, an engineering services firm in South Carolina, used OLCF leadership computing resources to develop add-on parts that will substantially reduce the fuel consumption of Class 8 long-haul trucks. Access to Jaguar enabled BMI to develop and use the most detailed, accurate numerical model of a Class 8 truck and trailer ever created. The aerodynamic results obtained when this model was used in conjunction with advanced NASA CFD codes have proven to be amazingly accurate against physical test data. And the rapid turnaround allowed a high degree of refinement of the parts in a short period of time. The first group of add-on parts is trailer aero improvements that can be easily retrofitted to existing or new trailers with little or no modification of the base vehicle. Class 8 trucks average 6 miles per gallon. Depending on the suite of parts chosen, long-haul trucks that are retrofitted can achieve an increase in fuel efficiency from 0.6 to 1 mile per gallon (equivalent to a 10% to 17% increase in fuel efficiency). This far exceeds the 5% improvement required by the California Air Resources Board for trucks operating in California. The result will be reduced fuel consumption by up to 3,700 gallons per truck per year and a reduction in CO_2 by up to 41 tons (82,000 lb) per truck per year.

The program is structured to address a range of HPC needs, including smaller, midsize, and large-scale allocations of computing time, under three participation categories.

- For smaller allocations (up to 2 million hours of compute time) companies can apply for a DD program allocation.

- For larger problems (up to 20 million hours) that are directly related to either the DOE energy mission or national emergencies or for projects that broaden the community of researchers capable of using leadership computing resources, companies can apply through the DOE ALCC program.
- Companies addressing the most complex scientific challenges (requiring more than 20 million hours and using 20% or more of Jaguar in their production runs) can apply for time through the INCITE program.

HPC has been called a game-changing technology. The OLCF HPC Industrial Partnerships Program is providing companies of all sizes and needs a gateway to this technology, and we are already seeing the results (Section 3).

SQUEEZING CO_2 TO THE MAX

Ramgen Power Systems, under sponsorship from DOE, is developing a unique compression product with a 100:1 compression ratio so that CO_2 gas can be pumped into storage at a significantly lower cost than with conventional carbon sequestration methods. The “Rampressor,” as it is known, generates shock waves and air compression in a manner similar to ramjet inlets on supersonic aircraft, hence the name. Access to Jaguar is significantly enhancing and accelerating Ramgen’s ability to conduct very high-resolution CFD simulations to characterize component, rotor, and stage performance and optimize the demonstration design to reach product-level performance. The Rampressor is projected to reduce the capital costs of CO_2 compression by 50% and produce a minimum of 25% operating cost savings. A complementary output of this project is improving the parallel implementation in the FINE/Turbo application from commercial software vendor NUMECA International. This benefits not only Ramgen, which has been using FINE/Turbo for prototype development, but also the broader industrial community of manufacturers that use NUMECA’s fluid dynamics software.

6. RISK MANAGEMENT

CHARGE QUESTION 6: Is the OLCF effectively managing risk?

The OLCF's Risk Management Plan (RMP) describes a regular, rigorous, proactive, and highly successful review process first implemented in October 2006. Operations and project meetings are held weekly, and risk is continually being assessed and monitored. A project/operation risk meeting is attended by the federal project director from the Oak Ridge DOE site office. The OLCF sends aggregated risk reports monthly to the DOE program office.

The OLCF has a highly successful history of anticipating, analyzing and rating, and retiring risk for both project- and operations-based risks. Our risk management approach uses the Project Management Institute's best practices as a model. The RMP includes

- identifying and analyzing potential risks,
- ensuring that risk issues are discovered and understood early on,
- ensuring that mitigation plans are prepared and implemented, and
- developing budgets with consideration of risk.

OLCF risk assessment is a six-step process. Once a risk is identified through a discussion of threats and vulnerabilities, the chance of occurrence is determined and its impact on project or operations scope, cost, and schedule are assessed. Then a (typically informal) cost/benefit analysis is performed to determine if mitigation activities are called for. If so a plan is made and executed when appropriate. Mitigation activities are reported and tracked as with any other project work breakdown structure (WBS) activity element, or if there are operational risks, they are reported and tracked as part of the periodic OLCF risk meetings.

Risk planning focuses on likelihoods and consequences. Likelihood is assigned as "very likely" (over 80%), "likely" (between 80% and 30%), and "unlikely" (below 30%). Impact category thresholds differ according to the impact area and whether the impact is to a particular project or to operations. For OLCF operations, the following table is used.

Category	**Impact on Project**		
	Low	**Medium**	**High**
Cost	<$250,000	>$250,000 & <$500,000	>$500,000
Schedule	<1 month	>1 month & <3 months	>3 months
Scope (based on performance metrics)	<10%	>10% & <20%	>20%
Other	Depends on the area of concern and is usually a subjective evaluation.		

A risk management software application helps the team to record, track, and report on identified project risks. The risk register software uses the assessment to rate and rank them as they are entered and updated over time. A risk rating is a dimensionless numeric score generated from a combination of likelihood and the highest rated impact, which is used to give a sense of relative importance.

At its periodic risk reviews, weekly staff meetings, and ad hoc discussions, the OLCF management team continues to focus attention on the high and moderate risks while keeping an eye on low risks, which may increase in importance over time. The managers and group leaders benefit from a thorough familiarity with previous risk profiles as they review the risk register, and they are in a strong position to anticipate future events. There were 173 risks registered for the OLCF-1 project that have been retired, and the OLCF-3 project team is collecting and assessing the risks associated with that new project.

Risk assessment is a major consideration for the DOE SC, and OLCF staff attended a Risk Management Techniques and Practice (RMTAP) workshop. The RMTAP workshop, which was sponsored by the ASCR program and the National Nuclear Security Administration (NNSA) Advanced Simulation and Computing (ASC) program, assessed current and emerging techniques for effectively identifying, understanding, managing, and mitigating the risks associated with acquiring leading-edge computing systems at high-performance computing centers (HPCCs). Representatives from 15 HPC organizations, four HPC vendor partners, and three government agencies attended. The top risk categories discussed there were system-scaling issues, requests for proposal/contract, acceptance testing, and vendor technical or business problems. The workshop concluded that HPC projects often require a tailoring of standard risk management practices and that the special relationship between the HPCCs and HPC vendors must be reflected in the risk-management strategy.

Several of the workshop best practices recommendations are standard OLCF practice, including

- developing a prioritized risk register with special attention to the top risks,
- establishing a practice of regular meetings and status updates with the platform partner,
- supporting regular and open reviews that engage the interests and expertise of a wide range of staff and stakeholders, and
- documenting and sharing the acquisition/build/deployment experience.

There are currently 25 entries 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 known risks in the OLCF Operations Risk Register.

Across-the-board

- Funding uncertainty is one of the highest risks for the OLCF. Annual budgets are set with guidance from the ASCR office, but actual allocated funds are unknown until Congress passes funding bills. Continuing resolutions are common, and we often go several months before actual funding is resolved. The risk is that we may have to delay some purchases, activities, hiring, etc., or possibly adjust lease payment schedules, resulting in substantially higher costs and perhaps schedule delays. We will continue to maintain close contact with the Federal Project Director and ASCR Program Office to understand the changing funding projections so that alternative plans can be made in sufficient time. Where possible, we will structure contracts to accommodate flexible payment terms. Rating: **High**

- That utility costs may escalate at rates higher than expected is another high risk for the OLCF. Utility costs may rise higher than the 3–4% annual increase that is budgeted. We will accept this risk, as there is little we can do to influence utility rates from TVA. However, we will monitor the political and economic environment closely and, if sufficient time exists, adjust operating budgets to accommodate any exaggerated increases. If this cannot be done, we may be forced to request additional funds from DOE, reduce some operations activities, or take the drastic step of powering down for some period of time. Rating: **Medium**
- Staffing is a concern. Much of the effort within the OLCF is provided by highly trained and highly experienced staff. The loss of critical skill sets or knowledge in certain technical and managerial areas may hinder ongoing progress. Good career development programs have been implemented within the division to retain high-quality personnel. Succession planning is promoted, and there are active laboratory-wide recruiting campaigns and outreach programs. Despite the best efforts in recruiting, training, etc., funding uncertainty continues to be a concern for the OLCF's ability to attract and keep the high-quality staff essential to its success. For example, several other risk register entries describe risk mitigation efforts involving SciComp, HPC Operations, and TechInt, whose contributions are critical to the mission of both the OLCF and DOE. Demands on these groups of specialists are increasing at an extraordinary rate and the danger remains that staff burnout will take its toll. Rating: **Medium**
- There is always a risk that the facility experiences a safety occurrence resulting in serious personal injury. We work to reduce these risks with monitoring of worker compliance with existing safety requirements, daily tool box safety meetings, periodic surveillances using independent safety professionals, joint walk-downs by management and work supervisors, and encouragement of stop-work authority of all personnel. Observations from safety walk-downs will be evaluated for unsatisfactory trends (e.g., recurring unsafe acts). Unsatisfactory performance will receive prompt management intervention commensurate with severity of the safety deficiencies. Integrated Safety Management is a core performance metric for the entire laboratory. Safety is a top UT-Battelle priority that carries throughout the laboratory, and the OLCF understands that it is critical to its success to provide a safe working environment. Rating: **Low**
- Because of its high electrical power use, the facility is at risk of fire. The OLCF follows and meets all electrical and fire protection codes, rules, and guidelines. Nonetheless, should fire or similar damage occur, the OLCF has implemented data integrity mechanisms and complete metadata backup so that recovery is possible and losses are minimized. Rating: **Low**
- System cyber security failures involving unauthorized access or use of systems may force a shutdown for extended periods or otherwise degrade system productivity. We have developed a cyber security plan that implements protection to the MEDIUM level. This includes such things as continual monitoring for security breaches, user identity checks prior to granting accounts, 2-factor authentication, and periodic formal tests and reviews. A U.S. government laboratory is subject to intense external assaults on its IT systems and networks. OLCF staff, in concert with ORNL's cyber security technical and policy teams, are constantly looking for ways to balance the protection of its IT resources with its need to continue its science mission. Rating: **Low**

System utilization

- Some users may determine that it is too much effort and/or expense to port their code to the new heterogeneous architecture that is planned for future systems. This will limit the usefulness and productivity of the new systems. The buy-in costs associated with new HPC technologies is almost always a significant entry or upgrade barrier. However, outreach, training, and the availability of libraries and development tools currently being investigated will ameliorate some of the possible resistance. Specifically, SciComp is analyzing six important target codes to provide a better understanding of the difficultly, cost, and time needed to upgrade to the planned heterogeneous architecture, the expected benefits of upgrading, and how best to proceed with porting efforts. As currently planned, the new architecture has sufficient base CPU power to provide most HPC user codes with significant performance increases without any special code modifications that take advantage of the accelerator chips. However, substantially better performance and, therefore, more benefit per unit of cost, would be obtained if the codes were modified to off-load some computations to the companion accelerator chips. An application readiness review is planned for August 2010 that will present SciComp's porting analysis findings. This will provide a much better understanding of the risk and the best response to those risks. A measure of success of the mitigation efforts will be how many users actually upgrade their code and to what extent. The details of this potential metric will be better understood after the pending Application Readiness Review. Rating: **Low**

- Related to the risk above is the situation where leadership-level computing is not achieved. Too many application runs may be submitted that do not achieve "leadership" status. The OLCF has established job queue policies with high preference for leadership jobs and continually evaluates their effectiveness. The OLCF is involved with the INCITE proposal selection process, which ensures that leadership projects receive allocation preference. SciComp has been established to help users scale their applications to leadership levels. Leadership computing is defined as utilizing a certain percentage of the available computing capability of a system. So far in CY 2010, Jaguar XT5 has been running about 44% of its jobs above 20%. Clearly the best chance for improvement in this area lies with SciComp helping scientists scale up their applications. Rating: **Medium**
- As new or upgraded computing platforms are acquired, applications may not be sufficiently prepared to take advantage of the increased computing capabilities. SciComp continues to work closely with the HPC user community to improve codes to take maximum advantage of any new technology that the OLCF introduces. Advanced test beds are frequently acquired to provide early access to new technologies. UAO and SciComp also conduct education, outreach, and training to continually expand and extend the skill levels of the HPC user base as well as ORNL staff. The OLCF has long recognized that our mission is to provide not only leadership computing systems, but also the support scientists need in order to use these systems effectively. Rating: **Medium**

Outages

- Power outages from external causes may create delays in user job completion or otherwise hinder system performance. The OLCF constantly evaluates risk in this area. It has installed cost-effective back-up capabilities (generators, UPS, dual-power cabinet designs, etc.). Cooling equipment failures are also possible. HPC systems operate with fairly strict temperature requirements. OLCF systems have automatic shutdown features in case temperatures rise above a set threshold. In addition, there are redundant chillers (five, where the systems could run on three). There are also redundant cooling towers and pumps, and it is possible to cross-connect the OLCF with the cooling capabilities of other facilities at ORNL. Rating: **Low**
- Network outages could prevent effective system use. If networks are inoperable or running too slowly, many users could lose access to the OLCF systems. There is some redundancy in ESNet, but there is significant residual risk there. ORNL has contracted with a commercial network

provider to supply alternative network capability, although that would be at reduced performance. Rating: **Low**

Performance

- Maintaining high availability and stability of systems is critical to users. There is a risk that the system stability and availability may not be sufficient to meet user needs. Policies have been implemented that control availability to minimize maintenance downtimes, coordinate upgrades, maximize fault-tolerant HW and SW, etc. Availability and stability are continually monitored in order to detect trends in time to take remedial action. Rating: **Low**
- Users require support (e.g., account management, help desk) to use large-scale computing systems effectively. There is a risk that the support we provide will not be adequate. To mitigate this risk, OLCF staff communicate frequently and directly with users, measure satisfaction with formal surveys, and use liaisons to get better insight into users' problems and issues. This risk is somewhat different from user dissatisfaction with system use due to technological inadequacies (e.g., poor system performance, unscheduled downtimes, lost data). Those are covered in other registry entries. This risk has to do with the interactions users have with UAO. Rating: **Low**

File systems

- Oracle has recently indicated that its support model for Lustre has changed. It has publicly announced that

  — Lustre will remain open source and released under the GPL,
  — Lustre 1.8 will continue to be supported through mid-2012,
  — Lustre 2 will be supported on both Oracle and non-Oracle hardware, and
  — Lustre 2 support on non-Oracle hardware will require certification.

  The OLCF is unlikely to upgrade Lustre 1.8 to Lustre 2 prior to early 2012 because, although Cray may receive special dispensation for continued support as they did with Lustre 1.6, moving the OLCF to Lustre 2 will require significant costs in upgrading our system to Oracle hardware or certifying our existing hardware configuration under Oracle's Lustre certification program. Since Lustre is open-source software, TechInt has increased the OLCF's engagement with the Lustre open-source developer community, hosting weekly Lustre community teleconferences with participation from a number of Lustre stakeholders, including Cray, DDN, LLNL, NASA, and SGI. Should Oracle reduce its support below acceptable levels, we will further increase our engagement with, and financial support to, the Lustre open-source developer community. In the interim we will work with Cray and Oracle on defining support options for Lustre 2 on the Spider parallel file system. Should certification be required, TechInt will work with Oracle and Cray to certify the Spider hardware configuration for Lustre 2. TechInt has established itself as a leader in the Lustre community, having developed deep expertise in Lustre development, configuration, deployment, and optimization. Our engagement with the Lustre community has fostered a number of collaborations that would provide the basis for a community-led Lustre support model. Rating: **High**
- Metadata performance is critical to a wide variety of leadership application workloads. Lustre is currently limited to a single metadata server, as Clustered Metadata has not been delivered. Single metadata server performance has been a bottleneck on occasion for a few of our application workloads. Metadata performance degradation has the potential for significant adverse impact both on applications and interactive users as the demand for system resources continues to increase. TechInt is working with Cray and Oracle to improve single-metadata-server

performance. This work includes SMP scalability and reduction in the overhead of distributed locking protocols. SciComp and TechInt will work with application teams to reduce their metadata workloads through code restructuring and the use of middleware I/O libraries. In addition, TechInt is developing tools to monitor and respond to metadata performance slowdowns in order to minimize the impact to the overall user population. Segmenting the storage system into multiple file systems to increase the number of metadata servers (one per file system) is also being evaluated. An additional metadata server is unaffordable, given current funding levels. If the mitigation efforts described are inadequate, we will have to accept that performance will be less than what it could have been. If a performance trend indicates harm to the science mission, we will seek additional funding. Rating: **High**

- OLCF storage systems could exhibit reliability and/or data integrity issues. Some problems develop over time or do not surface until particular workloads are encountered. We have architected multiple file systems, which are partitioned to contain failure domains. This reduces the impact of failures to a smaller subset of users by trading off the performance a single user could expect. TechInt continues to develop tools to help with thorough ongoing analyses of storage system components to identify single points of failure and double points of data loss. The results of these analyses will drive the design of parallel file systems and operational procedures that will isolate domains of failure and data loss, allowing quick response to problems. Tools that enable identification and isolation of data corruption (in the event of failure) are being developed as well. Procedures that minimize the probability of data corruption during failure are being developed, and operational staff are trained in these procedures. An e-mail list has been established that allows the DOE SC laboratories to share information regarding system failures, and operational staff from ORNL, LBNL, and ANL are sharing information regarding system failures in order to prepare and anticipate common pathologies. Ongoing vendor contract metrics were established to help limit our financial exposure, but these do not mitigate technical performance risks. The acceptance test included a 14-day stability test that should reduce the likelihood of encountering these issues after acceptance, and the use of RAID 6 storage systems with fault-tolerant features should minimize the possibility of data loss. However, none of the above can fully eliminate this risk. To facilitate timely response to suspected system events that could cause data loss/corruption, TechInt will continue to develop scalable tools to isolate failure domains and the impact to our users in a timely fashion. Should system failures manifest that could be mitigated by storage system enhancements (firmware updates) or file system enhancements, TechInt will work closely with DDN, Oracle, and Cray to implement these enhancements. Rating: **Low**
- Application and interactive use of the Spider center-wide file system may be adversely impacted due to contention between systems over the shared resource. To combat and mitigate this effect, TechInt works closely with Oracle to establish requirements and develop software to provide quality-of-service (QoS) levels for the OLCF's compute resources. Assigning QoS levels to each resource will allow operational staff to prioritize traffic from the interactive and primary compute platforms while avoiding starvation of other systems. This helps to ensure a responsive system for the user editing files and performing interactive work while avoiding significant impact to batch jobs. QoS for Lustre is based upon a set of experimental patches demonstrating the feasibility and performance benefits of using a network request scheduler (NRS). Initial work on the NRS involved developing a Lustre object storage simulator to enable the exploration of various scheduling strategies and their effect on performance. While this work was focused on scheduling the network requests to maximize sequential access to the disk and thereby improve the available bandwidth, the infrastructure developed is adaptable to the needs of prioritizing traffic between different clients and compute platforms, giving us the ability to set QoS levels. Testing at scale of an NRS prototype demonstrated improved performance over naive in-order request scheduling policies. Follow-on efforts include providing the ability to prefer requests from specific clients,

testing the ability of the NRS to provide QoS in an environment with multiple clusters of Lustre clients, and evaluating the overall impact of this approach on performance and software maintainability. Rating: **Low**

- Differences between Lustre versions on Spider and the Cray systems impedes integration. We will continue to coordinate our development activities with Cray and Oracle. We have tested 1.4 and 1.6, as well as 1.6 and 1.8, version interoperability with good results. Our current production system is based on 1.6 version clients and servers. Work on a function shipping mechanism (data virtualization service) can reduce this risk by allowing the center-wide file system to maintain a consistent Lustre version while I/O clients can access the storage via the function shipping mechanism. We are currently finalizing support from Cray with pass-through support from the Lustre development team. We plan to deploy a Lustre 1.8 parallel file system environment (servers) during 2010, which will require interoperability with our Lustre 1.6 clients on the Cray XT platforms. Recent testing has demonstrated the scalability and stability of this configuration, and we anticipate very few issues related to interoperability transitioning to this configuration. Once stabilized, all current Lustre 1.6 center-wide parallel file systems will be transitioned to Lustre 1.8 and non-Cray Lustre clients. Upon upgrade to the next major release of CLE, the Cray Lustre clients will be upgraded to Lustre 1.8 as well. Rating: **Low**

Development environments

- To use HPC effectively, a fully functional software development environment is necessary. The risk is that some of these tools may be inadequate to allow practical levels of productivity. An Application Performance Tools Group was created to own the problem. We surveyed users on their requirements in this area and the adequacy of the tools available or planned. Contracts were initiated with vendors to supplement the work of the Tools Group to obtain additional functionality.
  - **Compilers.** Platforms are changing rapidly, with increasing system heterogeneity as well as the requirement to extract unprecedented levels of parallelism from the applications. The commodity market is operating at a much lower scale and is not funding the development of compiler technology at the levels needed for HPC systems. The OLCF will track system requirements and compiler vendors and make targeted investments to meet specific OLCF needs. Additionally, the research community is being tracked for ways to bring needed capabilities into vendor-supported compiling systems. The OLCF participates in actions to develop a large HPC community that works in concert to remedy the situation.
  - **Debuggers.** On today’s large-scale systems, debugging support is limited, with only one debugger vendor (DDT) capable of providing debugging support at large scale (after our investment). As system scales continue to grow at a rapid pace, the scalability of debugging solutions needs to increase as well. In addition, high-performance analysis tools for inspecting data for the source of code errors is currently extremely inadequate. The OLCF will continue with targeted investment in improving debugging capabilities. Additionally, the research community is being tracked for ways to bring needed capabilities into vendor-supported debugging systems. The OLCF participates in actions to develop a large HPC community that works in concert to remediate the situation.
  - **Application performance tools.** Performance analysis is limited to runs of, at most, a few tens of thousands of cores. Our ability to understand application performance at the scales leadership applications are expected to run is extremely limited. The commodity market is operating at a much lower scale and is not funding the development of performance tool technology needed for HPC systems. The OLCF will continue with targeted investment in improving performance analysis capabilities. Additionally, the research community is being

tracked for ways to bring needed capabilities into vendor-supported debugging systems as the volume of data generated at large scale is large and new analysis techniques need to be developed. The OLCF participates in actions to develop a large HPC community that works in concert to remediate the situation.

Rating: **Medium**

7. CYBER SECURITY

CHARGE QUESTION 7: Does the OLCF have a valid cyber security plan and authority to operate?

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

The OLCF has the authority to operate for 3 years under a C&A plan approved by DOE on September 22, 2009. The OLCF is accredited at the moderate level of controls, which authorizes the facility to process sensitive, proprietary, and export-controlled data.

The Cyber Security effort at the OLCF has an active Cyber Security Program Plan (CSPP) in place and is proactive in maintaining it and in adapting to new challenges. DOE awarded the facility continued Authority To Operate on September 22, 2009, after an updated CSPP with enhancements in the user identity proofing process was submitted. The accreditation of ORNL cyber enclaves expires June 21, 2011. To continue to meet the requirements for operating as a moderate enclave, the plan is being further updated to reflect improved guidelines released by NIST Special Publication 800-53 rev. 3. A new CSPP document reflecting these changes will be submitted for review next year.

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

Department of Energy

Oak Ridge Office
P.O. Box 2001
Oak Ridge, Tennessee 37831

February 3, 2010

Mr. Mike Bartell, Director
Computing and Computational Sciences
Oak Ridge National Laboratory
UT-Battelle, LLC
Post Office Box 2008
Oak Ridge, Tennessee 37831-6404

Dear Mr. Bartell:

ACCREDITATION OF OAK RIDGE NATIONAL LABORATORY CYBER ENCLAVES-EXTENSION OF AUTHORITY TO OPERATE (ATO)

Reference: 1/21/2010 letter, Mike Bartell to Johnny O. Moore, subject same as above

An ATO was granted to six unclassified enclaves for a three-year period effective June 21, 2007. However, the Office of Science Program Cyber Security Plan that delineates cyber policies covering these and other enclaves is being revised and a reissue date is uncertain at this time.

Therefore, an extension of the existing ATO is granted for these enclaves to an expiration date of June 21, 2011.

If there are any questions or additional information is required, please contact Dave Rosine at 574-8640.

Sincerely,

Johnny O. Moore, Manager
ORNL Site Office
Designated Approving Authority

cc:
Qui Nguyen, AD-41, ORO
Douglas E. Alred, ORNL
Tina C. Heath, ORNL