Argonne NATIONAL LABORATORY

2016 OPERATIONAL ASSESSMENT REPORT

ARGONNE LEADERSHIP COMPUTING FACILITY

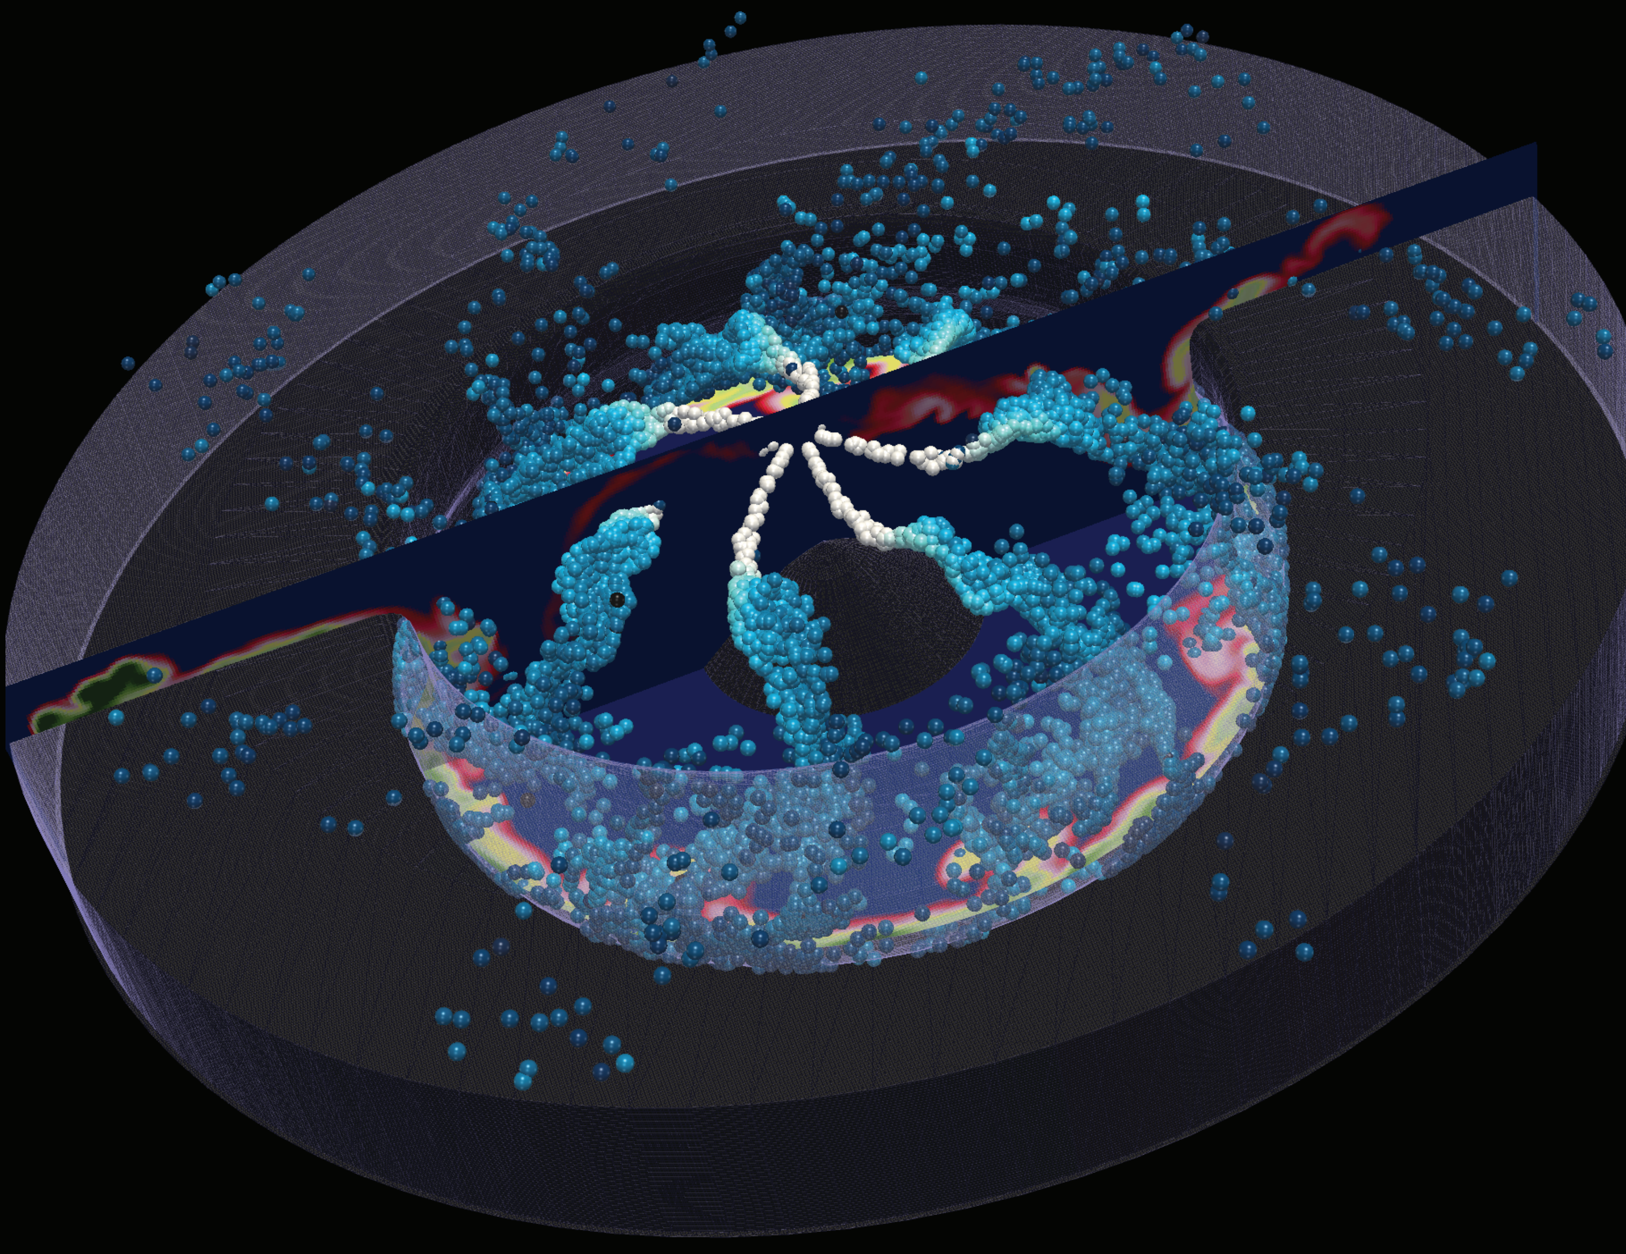

On the cover

This simulation captures fuel distribution inside the combustion chamber of a GM 1.9 L diesel engine being operated in an advanced low-temperature combustion mode. High-fidelity 3-D simulations provide unique insights into the fuel-air mixing process, which has enabled Argonne engineers to develop this combustion mode. It was produced with the computational fluid dynamics software CONVERGE and run on the Mira supercomputer at the Argonne Leadership Computing Facility at Argonne National Laboratory.

Image: Joseph A. Insley, Janardhan Kodavasal, and Sibendu Som, Argonne National Laboratory

Contents

Contents (Cont.)

Contents (Cont.)

Figures

Tables

This page intentionally left blank.

Executive Summary

In 2004, the U.S. Department of Energy's (DOE's) Advanced Scientific Computing Research (ASCR) program founded the Leadership Computing Facility (LCF) with a mission to provide the world's most advanced computational resources to the open science community. The LCF is a huge investment in the nation's scientific and technological future, inspired by a growing demand for capability computing and its impact on science and engineering.

The LCF operates two world-class centers, at Argonne National Laboratory and at Oak Ridge National Laboratory, and deploys diverse petascale machines that are 10 to 100 times more powerful than systems typically available for open scientific research. Strategically, the LCF ranks among the top U.S. major scientific facilities delivering impactful science. The work being done at these centers informs policy decisions and advances innovations in far-reaching topics such as energy assurance, ecological sustainability, and global security.

The leadership-class systems at Argonne and Oak Ridge run every day, 52 weeks a year. From an operational standpoint, the high level of service these centers provide and the exceptional science they produce justify their existence to the DOE Office of Science and the United States Congress.

This Operational Assessment Report describes how the Argonne Leadership Computing Facility (ALCF) met or exceeded every one of its goals as an advanced scientific computing center.

In 2016, the ALCF's primary resource, Mira, delivered 3.9 billion core-hours to 34 Innovative and Novel Computational Impact on Theory and Experiment (INCITE) projects and 1.7 billion core-hours to ASCR Leadership Computing Challenge (ALCC) projects (26 projects for the 2015–2016 ALCC year; 27 projects for the 2016–2017 ALCC year), and supported 180 Director's Discretionary projects. Furthermore, Mira had another excellent year in terms of overall availability (94.9 percent), scheduled availability (98.9 percent), and utilization (90.8 percent).

Moreover, ALCF's user community published more than 250 papers in high-quality, peer-reviewed journals and technical proceedings. The ALCF user experience team held numerous user workshops and tutorials on a wide range of topics and supported the laboratory's fourth highly successful Argonne Training Program on Extreme-Scale Computing.

As the LCF prepares to roll out powerful new supercomputers at both centers in the next couple of years (Argonne's future system, Aurora, is expected to deliver 200 petaflops), the technical staffs are already busy working to improve critical software and to develop tools that will work at these unprecedented scales. When the pre-exascale systems arrive, the Leadership Computing Facility will once again stand ready to deliver science on Day One.

Table ES.1 Summary of the Target and Actual Data for the Previous Year (2016) Metrics

Area	Metric	2016 Targets	2016 Actuals
User Results	User Survey – **Overall Satisfaction**	3.5/5.0	4.5/5.0
	User Survey – **User Support**	3.5/5.0	4.5/5.0
	User Survey – **Problem Resolution**	3.5/5.0	4.5/5.0
	User Survey – **Response Rate**	25.0%	45.1%
	% User Problems Addressed Within Three Working Days	80.0%	95.9%
Business Results	Mira Overall Availability	90.0%	94.9%
	Mira Scheduled Availability	90.0%	98.9%
	% of INCITE core-hours from jobs run on 16.7% or more of Mira (131,072–786,432 cores)	40.0%	76.4%
	% of INCITE core-hours from jobs run on 33.3% or more of Mira (262,144–786,432 cores)	10.0%	40.9%

Section 1. User Support Results

Are the processes for supporting the customers, resolving problems, and Outreach effective?

ALCF Response

The Argonne Leadership Computing Facility (ALCF) has processes in place to effectively support its customers, to resolve problems, and to conduct outreach. The 2016 user survey measured overall satisfaction, user support, and problem resolution, and thereby served both to mark progress and to identify areas for improvement (Table 1.1). The following sections describe ALCF events and processes; consider the effectiveness of those processes; and note the improvements that were made to those processes during calendar year (CY) 2016.

Table 1.1 All 2016 User Support Metrics and Results [a]

		2015 Actual	2016 Target	2016 Actual
Number Surveyed		990	N/A	1,019
Number of Respondents (Response Rate)		454 (45.9%)	25.0%	460 (45.1%)
Overall Satisfaction	Mean	4.5	3.5	4.5
	Variance	0.7	N/A	0.6
	Std Dev	0.8	N/A	0.7
Problem Resolution [b]	Mean	4.6	3.5	4.5
	Variance	0.5	N/A	0.4
	Std Dev	0.7	N/A	0.6
User Support [b]	Mean	4.5	3.5	4.5
	Variance	0.5	N/A	0.3
	Std Dev	0.7	N/A	0.6
		2015 Actual	**2016 Target**	**2016 Actual**
% User Problems Addressed within Three Working Days [c]		95.3%	80.0%	95.9%

[a] In September 2015, all Advanced Scientific Computing Research (ASCR) facilities adopted a new definition of a facility user based on guidance from the U.S. Department of Energy's (DOE's) Office of Science. Under this definition, a user must have logged in to an ALCF resource during a given time period. This definition of a user provides the basis for all survey results.

[b] Results in the table reflect the 2015 revised statistics for variance/standard deviation previously reported. The minor revisions had no material impact on statistical significance or conclusions.

[c] The statistical population represented in this metric includes problem tickets submitted from all users. Note that this is a larger population than that of DOE users.

Survey Approach

The ALCF contracted with Inquisium (Cvent's rebranded Web Surveys division) to manage the 2016 user survey. The team incorporated lessons learned from previous surveys and feedback from various ALCF teams, ALCF leadership, and the User Advisory Council (UAC), as well as DOE's Advanced Scientific Computing Research (ASCR) program.

1.1 User Support Metrics

In 2016, a total of 1,019 individual ALCF users met the definition of a facility user and were invited to complete a user survey. Of those users, 460 responded, for a 45.1 percent response rate — far superior to a generally accepted standard for survey response rates of 10% for this size of population. ALCF surpassed all targets for the survey metrics.

Table 1.2 displays responses grouped by allocation program. While Director's Discretionary (DD) and Innovative and Novel Computational Impact on Theory and Experiment (INCITE) users each reported higher average Overall Satisfaction than ASCR Leadership Computing Challenge (ALCC) users, the results are not statistically significant. Other metrics are comparable, in that the variations are statistically insignificant.

Table 1.2 2016 User Survey Results by Allocation Program

2016 Metrics by Program		INCITE	ALCC	INCITE + ALCC	DD	All
Number Surveyed		373	224	597	422	1,019
Number of Respondents		187	100	287	173	460
Response Rate		50.1%	44.6%	48.1%	41.0%	45.1%
Overall Satisfaction	Mean	4.6	4.4	4.5	4.4	4.5
	Variance	0.5	0.6	0.5	0.6	0.6
	Std Dev	0.7	0.8	0.7	0.8	0.7
Problem Resolution	Mean	4.6	4.5	4.6	4.4	4.5
	Variance	0.3	0.4	0.3	0.4	0.4
	Std Dev	0.5	0.6	0.6	0.6	0.6
User Support	Mean	4.6	4.6	4.6	4.5	4.5
	Variance	0.3	0.4	0.3	0.5	0.3
	Std Dev	0.5	0.7	0.5	0.7	0.6
All Questions	Mean	4.6	4.5	4.5	4.4	4.5
	Variance	0.3	0.4	0.3	0.4	0.4
	Std Dev	0.5	0.6	0.6	0.6	0.6

As Table 1.3 shows, in 2016, ALCF again exceeded the Overall Satisfaction and User Support targets.

Table 1.3 2015 and 2016 User Support Metrics

Survey Area	2015 Target	2015 Actual	2016 Target	2016 Actual
Overall Satisfaction Rating	3.5/5.0	4.5/5.0	3.5/5.0	4.5/5.0
Average of User Support Ratings	3.5/5.0	4.5/5.0	3.5/5.0	4.5/5.0

1.2 Problem Resolution Metrics

As shown in Table 1.4, ALCF exceeded the target set for the percentage of problem tickets addressed in three days or less. ALCF defines a ticket as "addressed" after the following conditions have been met: (1) a staff member has accepted the ticket; (2) the problem has been identified; (3) the user has been notified; and (4) the staff member has solved the problem or is actively working on it.

Table 1.4 Tickets Addressed Metric

	2015 Target	2015 Actual	2016 Target	2016 Actual
% User Problems Addressed within Three Working Days [a]	80.0%	95.3%	80.0%	95.9%
Average of Problem Resolution Ratings	3.5/5.0	4.6/5.0	3.5/5.0	4.5/5.0

[a] The statistical population represented in this metric includes problem tickets submitted from all users. Note that this is a larger population than that of DOE users.

1.3 User Support and Outreach

1.3.1 Tier 1 Support

1.3.1.1 Phone and E-mail Support

The ALCF answered and categorized 6,018 support tickets in 2016. As in past years, the largest number of tickets related to Accounts (see Table 1.5). There was a slight increase in tickets in this category from previous years, which can be attributed to tasks related to setting up Theta for early access. These tasks include setting up accounts for vendors and early access users and supporting development work on the next-generation architecture. In addition, work was being done to bring Theta into production. Also, ALCF staff noticed a drop in quota management tickets. This is likely due to users' increasing familiarity with the ALCF's use of the quota system.

Table 1.5 Ticket Categorization for 2015 and 2016

Category	2015	2016
Access	1,070 (18%)	1,058 (18%)
Accounts	2,532 (42%)	2,755 (46%)
Allocations	633 (11%)	603 (10%)
Applications Software	236 (4%)	179 (3%)
Automated E-mail Responses	553 (9%)	480 (8%)
Data Transfer	23 (0%)	41 (1%)
I/O and Storage	182 (3%)	184 (3%)
Miscellaneous	215 (4%)	212 (3%)
Quota Management	60 (1%)	27 (0%)
System	448 (8%)	467 (8%)
Visualization	10 (0%)	12 (0%)
TOTAL TICKETS	5,962 (100%)	6,018 (100%)

1.3.1.2 Theta ESP: From Touchpoints to Journeys

The ALCF employed a customer journey map approach for the Theta Early Science Program (ESP) to better understand the user's experience from initial contact, through the process of engagement, and into a relationship. Figure 1.1 shows a portion of a customer journey map for a typical customer identified generically as "Steve." ALCF staff looked at the user's entire arc of engagement across multiple touchpoints: person-to-person interactions and via the website, e-mail, phone, and mailing list. This approach will help set priorities for the most important gaps and opportunities to improve the user journey for early users on Aurora and the coming machines.

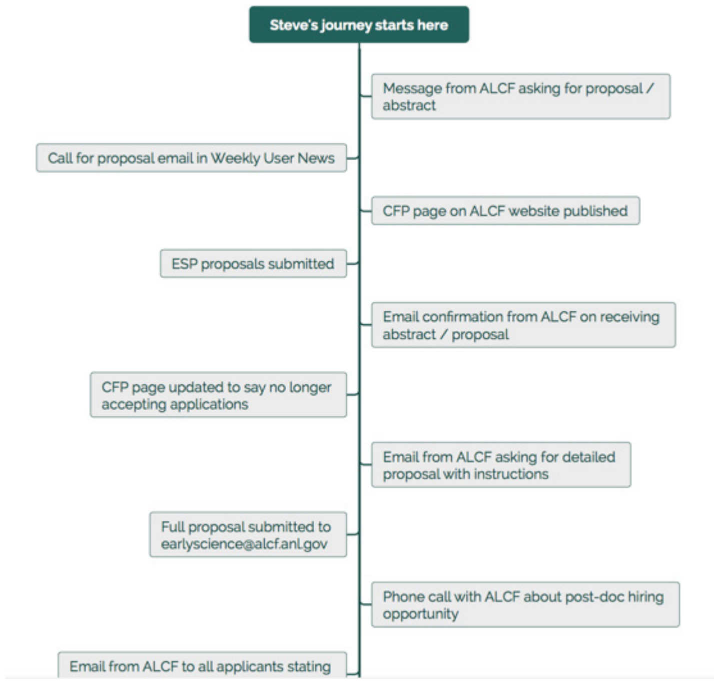

Figure 1.1 A portion of a customer journey map for "Steve."

1.3.1.3 User Onboarding

In 2016, the team adopted Agile methods for getting PIs, their projects, and their teams setup on ALCF systems—a process called onboarding. The method was used for ALCC and INCITE ramp-ups this year. The team used a combination of task list, document management, and process improvement software tools (Trello, Box, and OmniGraffle) to manage the process effectively.

The entire ramp-up process was divided into different phases ("Sprints" in Agile). This gave the team a framework for completing project work in a regular cadence and for identifying and mitigating any risks in meeting the objectives. To this end, for example, the team scheduled and ran daily Scrums, a feedback-driven approach where people meet in a room, quickly and transparently inspect a process, and determine next steps on a daily basis, to share the status of the deliverables with other team members. This practice gave the team members an opportunity to let each other know what their priorities were for the day, surface any potential roadblocks, and communicate any unexpected interruptions in completing their work. Overall, adopting Agile methods for these projects helped to increase the quality of the deliverables.

1.3.1.4 Continuous Improvement

While the ALCF has adopted more formal approaches to project and process management, the user experience team also practiced a continuous improvement approach to identify opportunities for streamlining processes. This team meets on a weekly basis to discuss opportunities for improvement and strategies for addressing them. Through these meetings, the team seeks ways to improve the processes and review how the implemented changes work for both users and the facility.

The process to request and obtain an ALCF account has been streamlined in the following ways: (1) the names of the fields and the text in the user-facing forms have been edited to provide more clarity; and (2) a drop-down selection of alphabetically listed countries has been added in place of a free-form text field, allowing users to select an option at a glance. These changes have reduced the number of follow-up e-mail exchanges needed to obtain a complete set of information. Changes like these will help the facility to have accurate and consistent data stored in our database for users.

1.3.1.5 Processing User Accounts for Workshops

ALCF staff members assisted with various workshops and webinars aimed at training and supporting the users and their project teams. For these events, staff members set up user accounts, handled project set-up, and helped facilitate the registrants' ability to participate.

The ALCF hosted numerous workshops and webinars in 2016, either onsite at the laboratory or offsite as a primary organizer or contributor. These included the following events:

- 2017 INCITE Proposal Writing (webinar)
- Scaling Your Science on Mira
- Ensemble Jobs for Better Throughput
- Argonne Training Program on Extreme-Scale Computing (ATPESC)
- Virtual Engine Research Institute and Fuels Initiative (VERIFI) Workshop
- QMC Training Workshop
- Theta Early Science Program Workshop
- Intel Xeon Phi Users Group (IXPUG) 2016 Annual Meeting
- Getting Started on ALCF Resources (multiple webinars)
- Cray Urika-GX Training Workshop

1.3.2 PI Support

1.3.2.1 Scheduling Decision Support

The ALCF makes scheduling decisions and adjustments throughout the year to help INCITE and ALCC project teams achieve their allocation goals on Mira. The ALCF publishes a graph of the core hours used, by program, per week. This graph in Figure 1.2 presents several data points in a single image to show how scheduling decisions affect the usage throughout the year. It is updated on a daily basis and available on the ALCF's reporting site for ALCF staff.

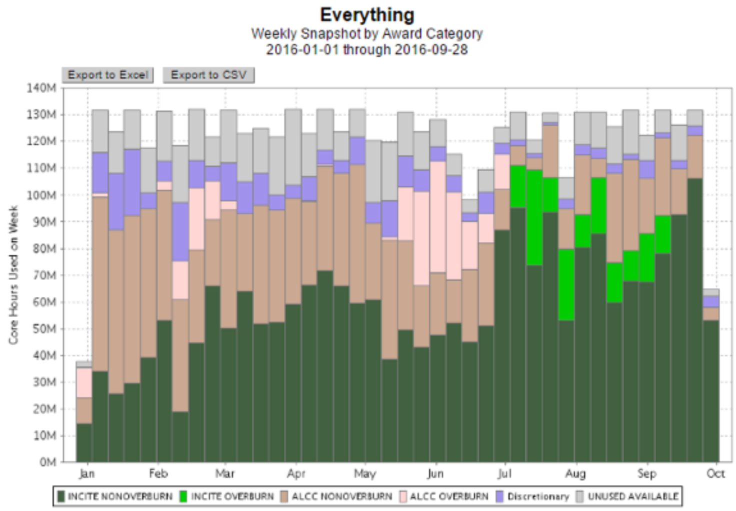

Current Week totals. Note: The end date for the week is 2 days ago due to jobs spanning 2 calendar days.

START DATE: 2016-09-25 END DATE: 2016-09-28

	Core Hours	Percent of all
Incite % Machine Needed		60.0%
Incite Total	53,172,861	70.4%
Incite Overburn Only	0	0.0%
ALCC Total	4,804,724	6.4%
ALCC Overburn Only	0	0.0%
Discretionary	4,143,666	5.5%
Unused Available	2,55,21,48	3.4%
Machine Unavailable	10,824,073	14.3%

Figure 1.2 Snapshot of a daily graph used by schedulers to make adjustments on Mira. The data table below the graph shows the current week's information. Download links give the viewer all of the graph data for use in other applications. The brighter colors indicate the results of applying the ALCF policy changes.

1.3.3 Application Support

1.3.3.1 Individual Projects

QMCPACK Optimizations for BG/Q

QMCPACK is a mature, open-source code capable of performing many-body electronic structure calculations for general materials and molecules. This high-performance code employs a modular, object-oriented design in C++, with extensive use of generic, template-based programming to maximize code reuse for operations such as mixed-precision calculation. For scalable parallel execution, QMCPACK employs a hybrid MPI+X programming model, with X being OpenMP on symmetric multi-processing (SMP)/many-core machines and compute unified device architecture (CUDA) on graphics processing units (GPUs). QMCPACK is being developed in collaboration with Oak Ridge National Laboratory, Sandia National Laboratories, and Lawrence Livermore National Laboratory and has a joint INCITE allocation (PI: David Ceperley) at ALCF and the Oak Ridge Leadership Computing Facility (OLCF). By rewriting the most computationally intensive parts of QMCPACK using compiler-specific extensions (called vector intrinsics) to better utilize the IBM Blue Gene/Q (BG/Q) processor, ALCF staff members Anouar Benali, Ye Luo, and Vitali Morozov were able to improve QMCPACK performance on Mira by

30 percent. In addition, by rewriting code to use single-precision instead of double-precision in key data structures, they decreased the amount of data that needed to be stored in memory by 45 percent.

Scaling RS3D on Mira

In support of a multiyear Director's Discretionary allocation (PI: Yun-Xing Wang), ALCF staff member Wei Jiang worked on a major rewrite of the RS3D code. This code uses small-angle X-ray scattering (SAXS) data and known secondary structures as input. Starting from a glob model at a nucleotide level of resolution, the algorithm carries out natural hierarchical moves based on the structural composition of ribonucleic acid (RNA). Three major improvements were performed: (1) parallelization of the original code using a message passing interface (MPI); (2) optimization of an existing OpenMP implementation; and (3) use of quad processing extension (QPX) intrinsics in computationally intensive kernels. The new version of RS3D is more than 4× faster than the original serial code. This work was performed under a Director's Discretionary project. The team submitted a 2016 INCITE proposal.

Improving Earthquake Simulation Workflows

In support of an INCITE 2016 allocation (PI: Thomas Jordan), ALCF staff member Tom Uram wrote a Python module to improve the project team's ability to use Mira at capability scale. The team uses Python to generate seismic events, which are then simulated using SORD, a rupture dynamics modeling code. The Python module is a more integrated and seamless approach to the workflow and has enabled Jordan's team to run 16-rack and 32-rack jobs.

Aytekin Gel Project

In support of an ALCC 2015–2016 allocation (PI: Aytekin Gel), ALCF staff member Ray Loy reworked workflow software to allow the creation of capability-sized ensembles of varying-sized multiphase flow with interphase exchanges (MFiX) jobs. MFiX, developed at the National Energy Technology Laboratory (NETL), simulates the hydrodynamics, heat transfer, and chemical reactions in fluid-solid systems. ALCF staff member William Scullin identified and patched issues with MFiX's memory allocation that had prevented the use of four or more ranks per node on Blue Gene/Q systems. He also resolved minor issues in code style to allow the compilation of MFiX with IBM's vectorizing compilers. Benchmark and scaling data collected by Scullin identified a bottleneck at 4,096 MPI-ranks beyond which Cartesian cell cutting routines cannot scale.

Accelerating Density Functional Theory Calculations with Advanced Algorithms

In support of an ALCC 2015–2016 allocation (PI: Robert DiStasio, Jr.), ALCF staff member Alvaro Vazquez-Mayagoitia worked with the project team on several methods for accelerating the evaluation of the exact-exchange in Quantum Espresso. He also proposed a new solver for the Poisson equation with a smooth, preconditioned, and conjugated gradient for rapid evaluations of exact-exchange that is highly efficient for many-core processors. The results were presented at Intel Xeon Phi Users Group (IXPUG) 2016 and the 2016 Intel HPC Developer Conference. A publication is in preparation. These advances were realized with the help of a summer student from Cornell University.

Optimization of Many-body Water Force Field

In support of a Director's Discretionary allocation (PI: Francesco Paesani), ALCF staff member Chris Knight optimized a many-body with polarization model for water. Originally, computation of the many-body interactions accounted for 80% of runtime on 2,048 BG/Q nodes for a system with 17,728 water molecules. Moreover, this system did not scale with regard to MPI; the code was not multi-threaded; and it had an unnecessarily large memory footprint, which prevented running with many MPI ranks per node. Knight's efforts resulted in an 8–10× speedup of the same system running on 64–512 BG/Q nodes, a 100× memory reduction, and good parallel scaling of many-body interactions with regard to both MPI and OpenMP. The team submitted a 2017 INCITE proposal, and it has one manuscript in review.

Tuning Hybrid Monte Carlo for Gauge Field Evolution

In support of an ALCC 2016–2017 allocation (PI: George Fleming), ALCF staff member James Osborn implemented a force-gradient update into the molecular dynamics (MD) evolution step for the gauge field, which eliminated the 3rd-order errors in the time step. Along with ALCF postdoctoral researcher Xiao-Yong Jin, they have been tuning the mass preconditioning parameters with the newly included force-gradient step. The newly tuned configuration generation code takes about 40% less time to run than before. The same method should be applicable to the MIMD lattice computation (MILC) configuration generation runs under the INCITE project (PI: Paul Mackenzie), on which they plan to test it next.

Resolving HDF5 Issues on Mira File System

In support of an INCITE 2016 allocation (PI: Greg Werner), ALCF staff member Ray Loy worked with HDF5 developers to fix an issue in HDF5 that was blocking the project from running on Mira. The project required HDF5 features in the recent 1.10.0 release. Building this version on BG/Q revealed a code portability issue that Ray was able to fix. The change was reported to the developers and included in the next patch release. However, once built, the project team encountered problems with file locking errors related to Mira's filesystem configuration (but not exclusive to BG/Q). An additional patch was developed and installed to resolve the issue.

1.3.3.2 Preparations for Knights Landing and Theta

FHI-aims/ELSI Optimizations

In support of a Tier 2 Theta Early Science Program allocation (PI: Volker Blum), ALCF staff member Alvaro Vazquez-Mayagoitia collaborated with Blum's team to improve the performance of FHI-aims, an all-electron, full-potential electronic structure code, and the parallel electronic structure infrastructure library (ELSI) for the Knights Landing (KNL) architecture. ELSI interfaces to the Eigenvalue SoLvers for Petaflop-Applications (ELPA) and other massive parallel eigensolvers, and has been highly optimized (with the help of Intel performance engineers) to obtain specific kernels for fast QR decomposition with AVX512 intrinsics. Currently, it is possible to solve a dense eigenproblem for the full spectrum for double-precision symmetric matrices of order one million MPI ranks in less than 2.5 hours using 3,072 nodes of Theta (196,608 cores) — an unprecedented benchmark in petascale computing. This progress was presented at a 2016 FHI-aims developers meeting and at the Technical University of Munich. A comprehensive analysis of FHI-aims performance in Theta will be presented in two talks at the 2017 American Physical Society March Meeting. A technical report is in preparation.

QMCPACK Optimizations

In collaboration with ALCF's Intel Parallel Computing Center (IPCC), staff members Anouar Benali and Ye Luo, along with Intel engineers, worked on three important kernels for the QMCPACK code in preparation for Theta. By moving code from double precision to single precision, they achieved an overall speedup of 1.6× on KNL and a speedup of 1.2× on Blue Gene/Q (gained from memory bandwidth). Other kernels were encapsulated in a mini-app for performance optimization. Kernels were rewritten using nested threading to improve the performance per node and stencil methods for improved use of caching. These last optimizations are being ported to the code, and a paper with the detailed results has been accepted to the 2017 IEEE International Parallel & Distributed Processing Symposium.

GAMESS Optimizations

In support of a Tier 2 Theta Early Science Program allocation (PI: Mark Gordon), ALCF staff members Yuri Alexeev, Graham Fletcher, and Spencer Pruitt, and Intel engineers, collaborated on improving the performance of the General Atomic and Molecular Electronic Structure System (GAMESS), a general *ab initio* quantum chemistry package. The team threaded the self-consistent field (SCF) in the Hartree-Fock (HF) method; made integral code thread safe; and rewrote code in HF to alleviate the impact of the irregular access patterns and computation imbalance. As a result, the team achieved a speedup of 4–6× on a single Xeon PHI for large systems. The code restructuring decreased the overall memory footprint by more than 2× and also paved the way for integrating vectorized integral packages into the GAMESS codebase. The results of this work were presented at 2016 IXPUG and submitted to the SC16 technical program. There is a plan to integrate the ALCF-developed distributed data interface (DDI) library in the official GAMESS release in 2017. The key feature of this library is elimination of data servers by using an MP3 nonblocking mechanism with Casper. (Casper is a process-based asynchronous progress model for MPI remote memory access (RMA) communication on multicore and many-core architectures developed by Argonne's Mathematics and Computer Science [MCS] division.) As a result, a significant gain in performance and a reduction of memory footprint in GAMESS is expected. This work leveraged the expertise of two IPCCs, one at Lomonosov Moscow State University and the other at Iowa State University.

1.3.4 Resource Support

1.3.4.1 General Support

Cobalt on the XC 40

Cobalt now has the ability to interface with and run Cray systems starting with Cray's XC 40 platform. This capability has allowed the ALCF to drive our successful acceptance testing of Theta using our scheduler, and significantly broadens the types of systems onto which we may install Cobalt. This functionality still leverages much of the functionality that had been developed for the BlueGene platform as well as our data-analytics clusters, but leverages a new internal architecture in the system interface that should enable much more rapid porting of features developed for Cray platforms to general cluster platforms, where appropriate.

GHI: Hierarchical Storage Management and Backup Solution

GHI supplies a hierarchical storage management and backup solution for the general parallel file system (GPFS) file systems called the GPFS HPSS (high-performance storage system) Interface (GHI). It is fully automated and works quietly in the background providing space management and disaster recovery. In July of 2016, GHI was deployed on one of our project file systems with the goal of freeing up space occupied by unreferenced data. ALCF initially targeted expired project data by migrating it to the HPSS archive using GHI.

Once migrated, the data is considered “managed” by GHI and may optionally be purged from the file system to free space. Ultimately, all file system data will be migrated to the HPSS archive where a GHI backup will be run to collect and store the metadata for the entire file system, and it will be run at regular intervals going forward. The backup images will provide disaster recovery in the event of file system loss. The GHI restore process restores the file system namespace, extended attributes, and GPFS file metadata only. This selection allows recovery to be accomplished in a reasonable timeframe. The data portion of files that will be referenced immediately can be recalled in bulk through GPFS policies.

Improved Allocation Accounting Tool (sbank)

sbank is an improved allocation accounting tool. It replaces cbank, the tool written when the ALCF started. In comparison to cbank, sbank provides much more flexibility in querying its data. Users can control what fields are displayed, as well as the order and width of each column. Time range querying is easier as any time format will be accepted that is not ambiguous. Internally, the system is completely transaction based. No data are edited or deleted; rather, there are compensating transactions applied. This approach allows detailed history of the allocation. There is also a clean Python application program interface (API) that can access data. This tool will be used when ALCF starts performing real-time allocation enforcement in Cobalt (currently, allocation data are updated hourly, and Cobalt checks at job submission only).

1.3.5 Outreach Efforts

1.3.5.1 General Outreach

User Advisory Council

The ALCF’s User Advisory Council meets on a monthly basis and provides valuable input on key technology upgrades, messaging and communication, and user-centric metrics. In 2016, the seven-member UAC represented all three allocation programs (INCITE, ALCC, and Director’s Discretionary). In 2016, UAC members evaluated and provided feedback on the following ALCF activities: the user support training calendar; communications related to the ALCF Data Science Program and cross-DOE facility training; the system accounting software beta program (sbank); and software support and management decisions for Mira. The UAC also served as a sounding board for communication strategies related to DOE user data requests. Finally, council members reviewed the 2016 user survey and provided a thorough end-to-end beta test.

Support of Brewer Science

Brewer Science is a small business in Rolla, Missouri, that delivers innovative material, process, and equipment solutions for lithography, advanced packaging, microelectromechanical systems

(MEMS), nanotechnology, optoelectronics, and compound semiconductor applications. Brewer Science has some in-house computational science expertise, but no HPC systems and little knowledge about appropriate software and scaling. ALCF has established a partnership among Brewer Science, ALCF, Argonne's Laboratory Computing Resource Center (LCRC), Argonne's Materials Sciences Division (MSD), and the University of Chicago's Institute for Molecular Engineering (IME), and acts as project coordinator. LCRC is providing the initial computing time as Brewer Science learns and scales up. MSD and IME are providing the software expertise. ALCF will help with scaling and provide leadership-class resources as the projects expand. Brewer has successfully completed one project under this model and has three more in various states of progress.

ALCC Industry Agreement

In 2016, DOE authorized a special user agreement for ALCC projects from industry. The ALCF worked with Argonne's legal department, the DOE-Argonne Site Office (DOE-ASO), and DOE ASCR to put in place an agreement available to all ALCF industry partners. The agreement consisted of the Argonne Proprietary Master User Agreement and a rider that waives cost recovery as long as specific results are identified to be published by the company.

Industry Lunch at SC Conference

ALCF organizes a lunch for industry users of HPC at the International Conference for High Performance Computing, Networking, Storage, and Analysis, commonly referred to as SC. This lunch is an opportunity for the ALCF Division Director and ALCF leadership to discuss with industry leaders how to better work together. The industry invitees are a mix of current and potential ALCF users and have provided invaluable feedback. For example, at SC16, the need for training around HPC and DOE resources was brought up by almost all of the industry people. Attendees were from a wide range of businesses: General Electric, John Deere, Tri-Alpha Energy, United Technologies, Disney Animation Studios, and Dow Chemical.

HPC Industry Showcase

ALCF's Manager of Industry Partnerships and Outreach, David Martin, co-chaired the HPC Impact Showcase at SC16. The purpose of the Showcase is to show the real-world impact of HPC on industry products and practices. The format of the Showcase is a series of invited presentations from leading HPC practitioners in industry. The Showcase was the most successful one to date. The eleven talks over three days received overwhelmingly positive reviews, and average attendance was over 100, with a high of 150.

NDA Management

While it is the policy of ALCF to not share private project information with others, many academic and industry users seek specific legal protection of their proprietary information (such as geometries, code, or techniques) that may be exposed in the course of using ALCF systems or through their interactions with ALCF staff members. In such cases, ALCF works with Argonne's legal department to establish a nondisclosure agreement (NDA) that both meets the needs of the individual user and satisfies Argonne's directives from DOE. ALCF maintains a repository of these agreements for staff to review when working with outside organizations. Whenever dealing with vendors, ALCF actively seeks NDAs and follows a similar process for putting them in place. In some cases, ALCF extends these agreements to other organizations so that confidential information can be shared among the groups. Having an established procedure and repository has enabled rapid discussions with potential vendors and has simplified communications of NDA restrictions within ALCF and the Computing, Environment, and Life Sciences (CELS) directorate.

DOE Exascale Requirements Reviews

ALCF, in collaboration with ESnet, the National Energy Research Scientific Computing Center (NERSC), and OLCF, held a series of Exascale Requirements Reviews to determine the mission-critical computational science objectives through 2025 for each of the six DOE Office of Science program offices – High Energy Physics (HEP), Basic Energy Science (BES), Fusion Energy Science (FES), Biology and Environmental Research (BER), Nuclear Physics (NP), and Advanced Scientific Computing Research (ASCR). These workshops brought together key domain scientists and computational science experts to identify the requirements for developing an exascale ecosystem that includes computation, data analysis, software, workflows, HPC services, and other features. The reviews for HEP and BES were held in 2015. As of September 2016, this ASCR facility collaboration completed the remaining reviews: FES, BER, NP, and ASCR. Besides managing the FES and ASCR reviews, ALCF is developing the reports for all six of the reviews. ALCF also built and will be hosting the archival website for all materials developed from the reviews.

1.3.5.2 Workshops and Webinars

The ALCF conducted workshops and webinars to support the efforts of users and their project teams. ALCF also collaborates with peer DOE institutions to develop training opportunities, explore key technologies, and share best practices that improve the user experience.

Getting Started Videoconferences

To better meet user needs, the training schedule was expanded from 12 sessions in 2015 to 14 sessions in 2016. The sessions were also distributed more evenly throughout the year to allow users to attend training at the time it was most relevant to them. This approach netted a 16% increase in students over the previous year.

2017 INCITE Proposal Writing Webinar

To prepare prospective and returning users for the annual open call for proposals, the INCITE program office conducts INCITE proposal writing webinars. Free and open to the public, these webinars are conducted by program office staff and representatives from ALCF and OLCF, to provide guidance for writing an effective INCITE proposal. These webinars took place on April 13, 2016, and May 19, 2016, and hosted 27 and 47 participants, respectively.

Scaling Your Science on Mira

The eighth annual ALCF HPC machines scaling event drew new and experienced supercomputer users from around the globe. This intensive, 3-day event is the cornerstone of ALCF's user outreach program and gives prospective INCITE users the opportunity to work with ALCF staff and other experts, including tool and debugger vendors. The bulk of this year's event was devoted to hands-on application tuning. In addition, ALCF staff gave talks on a range of topics, including ensemble jobs, parallel I/O, and data analysis. Of the 64 attendees, nine teams or individuals stated a pre-event goal of submitting a new 2017 INCITE project proposal; of these, four did submit a proposal, and three others later applied for and received ALCC projects. Two teams or individuals stated a pre-event goal of submitting a renewal project for 2017; of these, one submitted a renewal and one submitted a new project proposal.

Ensemble Jobs Videoconferences

First offered in July 2015 and September 2015, the Ensemble Jobs for Better Throughput videoconference is for users whose workloads include small jobs (<8K nodes) that are able to run concurrently. ALCF staff showed participants how to set up an ensemble job script and to understand which job submission type is best for their project. A March session was added to the 2016 schedule. Overall enrollment in this popular webinar increased 28% over 2015.

ATPESC 2016

The ALCF ran its fourth successful Argonne Training Program on Extreme-Scale Computing (ATPESC) program last July and August. ATPESC, founded in 2013, is an intensive, 2-week summer program organized around seven core program tracks and focused on HPC methodologies that are applicable to both current and future machine architectures, including exascale systems. More than 60 postdoctoral students and early career researchers attended the program consisting of technical lectures, hands-on exercises, and dinner talks. To extend the reach of this program, ALCF staff produce and upload video playlists of ATPESC presentations to Argonne's YouTube training channel. Eighty-two talks from the 2016 training course were uploaded in September, and promotion is ongoing. Since the program's inception in 2013, ATPESC videos have been viewed more than 18,000 times. ATPESC organizers work diligently to improve the program from year to year. In 2016, this effort included surveying participants for feedback at the conclusion of each of the seven program tracks, as well as administering an overall program assessment survey at the conclusion of the event. Data collected from the surveys is shared with program track leads and mined for improvement ideas.

Lattice for Beyond the Standard Model Physics 2016 Workshop

On April 21–22, Argonne hosted the "Lattice for Beyond the Standard Model Physics 2016" workshop, co-sponsored by ALCF and Lawrence Livermore National Laboratory. The workshop brought together experimentalists, lattice gauge theorists, and continuum phenomenologists to discuss the role that numerical simulations can play in searches for new physics beyond the Standard Model.

Collaboration: VERIFI Workshop

The Virtual Engine Research Institute and Fuels Initiative (VERIFI) is an Argonne-led program focusing on improving internal combustion engine performance though advanced simulation and experimentation. ALCF staff teamed up with VERIFI scientists and researchers to organize a hands-on session as part of the 2016 VERIFI Workshop, held June 22–23, 2016, to showcase current work and to attract new industry partners. As part of the hands-on session, 14 participants received access to Mira, along with step-by-step instructions on job submission, details on Global Sensitivity Analysis for engine simulations, and guidelines for using the Swift parallel scripting language to launch thousands of simultaneous jobs of the CONVERGE application, the Computational Fluid Dynamics software from Convergent Science, Inc. Each participant ran 64 piston engine simulations on a total of 8,192 cores to see the benefits of using HPC for understanding and accelerating piston engine studies.

IXPUG

Argonne is a charter member and has representation on the steering committee of the Intel Xeon Phi Users Group (IXPUG), a community-driven organization that fosters technical collaboration around tuning for the Intel Xeon Phi coprocessor family. ALCF hosted the IXPUG annual meeting at Argonne, September 19–22, 2016. The meeting attracted 141 people who registered from 36 different institutions from as far away as India and Russia. The program included 8 Tutorials, 4 Keynotes, 27 Presentations, and 17 Lightning Talks. Feedback on the event was overwhelmingly positive. In addition, many of the presentations and talks were recorded by Argonne and placed on the IXPUG website.

ALCF Theta ESP Hands-On Workshop

On August 16–19, 2016, the ALCF hosted researchers from its Theta Early Science Program (ESP) for a workshop to help them port, benchmark, and optimize their applications on the facility's next-generation Intel Cray system. The ESP teams worked directly with ALCF staff and representatives from Intel and Cray so their codes could be run on the actual machine.

ALCF, NERSC, and OLCF HPC Software Webinar Series

ALCF, OLCF, NERSC, and the Interoperable Design of Extreme-scale Application Software (IDEAS) project launched a series of seven webinars called "Best Practices for HPC Software Developers" to help users of HPC systems carry out their software development more productively. For the first session on July 14, 2016, ALCF Director of Science Katherine Riley presented on "How the HPC Environment is Different from the Desktop (and Why)." Four of the presenters in the series are from the IDEAS project. The other three are from the national scientific user facilities.

1.3.5.3 Community Outreach

ALCF staff are extending the facility's outreach to cultivate and educate the HPC workforce of the future — elementary through undergraduate students — who will be driving future discoveries in the exascale age. These efforts have taken off in the past six years and are largely undertaken by staff members who volunteer their time.

Youth Outreach

Women in STEM (Science, Technology, Engineering, and Math)

Every February, Argonne hosts Introduce a Girl to Engineering Day for local eighth grade girls and pairs them with Argonne mentors to participate in group activities focused on STEM careers. ALCF staff member Jini Ramprakash serves on the organizing committee, and many other ALCF volunteers serve as mentors or as activities supervisors.

Other efforts that ALCF contributes to, or participates in, that promote STEM careers to girls and women include Women in Science and Technology (WIST), Argonne's annual Science

Careers in Search of Women conference, the Anita Borg Institute, the Grace Hopper Celebration of Women in Computing, and the Rocky Mountain Celebration of Women in Computing.

Argonne/ACT-SO High School Research Program

The Afro-Academic, Cultural, Technological, and Scientific Olympics (ACT-SO) is a youth initiative of the National Association for the Advancement of Colored People (NAACP) to promote academic and cultural excellence among African-American high school students. The Argonne/ACT-SO High School Research Program, now in its third year, provides mentors and facilities to help students prepare their research for competition. The Argonne volunteer mentors work with the students over a period of seven months. Last year, ALCF User Experience Analyst Liza Booker and her student participated in a local competition at the College of DuPage, where they won the silver medal in computer science.

Hour of Code

As part of Code.org's annual event, ALCF staff members give talks, demos, and lead Hour of Code activities in various Chicago and suburban classrooms during Computer Science Education Week (CSEdWeek). CSEdWeek, established by Congress in 2009 to raise awareness about the need to elevate computer science education at all levels, is observed each year during the week of computer scientist and U.S. Navy Admiral Grace Hopper's birthday on December 9.

Summer Coding Camp

In 2015, ALCF teamed with Argonne's education office to launch an intensive Python coding camp for high school students in conjunction with the DuPage ACT-SO program and Downers Grove public schools. The four-day course, developed and delivered by ALCF computer scientists, emphasized problem-solving skills and hands-on coding activities. In summer 2016, the ALCF team and Argonne's education office continued the program, expanding the course to five days and adding a section on robotics.

Computer Science Curricula

Several years ago, ALCF Deputy Director of Operations Ti Leggett initiated, developed, and taught a multi-week computer programming course at the Oswego Public Library and, later, the Yorkville Public Library. This curriculum has since been repurposed and adapted numerous times and for numerous audiences, ranging from high school teachers to undergraduate coding camps to 4-H Clubs. The original curriculum, which was aimed at 7- to 14-year-olds, is built around computer programming exercises from code.org and the popular game Minecraft.

University Outreach

Computational Science Coursework

ALCF computational scientist Christopher Knight played a large role in organizing and executing the Harvard University graduate course Extreme Computing: Project-based High Performance Distributed and Parallel Systems. This applied computation course, part of a category of courses developed by Harvard's Institute of Applied Computational Science, highlights the use of computers to solve scientific problems and allows students to work alongside practicing experts.

Skyway STEM Community College

The Skyway Conference is a group of eight regional community colleges that compete in sports and cooperate on initiatives in the arts and sciences. The Skyway STEM Poster Competition is an option for students in their first two years of college to explore their interests in STEM. Since 2012, volunteers from Argonne's postdoc community have served as the exclusive judging team for the yearly competition. Last year, ALCF postdoctoral researcher Preeti Malakar helped judge the competition at Waubonsee Community College.

Summer Student Program

Every year, the ALCF solicits project proposals from staff members who are interested in mentoring summer students. Through programs like DOE's Science Undergraduate Laboratory Internship (SULI) program and Argonne's Research Aide Appointments, college students are brought in to work alongside ALCF mentors on real-world research projects. In 2016, the facility welcomed its largest summer student class to date—a group of 24 students who ranged from college freshmen to PhD candidates. The students worked on a wide variety of projects that covered everything from system administration and data analytics to computational science and performance engineering.

Community Building

National Lab Science Day

On April 20, 2016, David Martin and Paul Messina attended the National Lab Science Day hosted by Energy Secretary Ernest Moniz in Washington, D.C. Martin contributed to the Computation exhibit, one of five exhibit areas set up for the event, and spoke with several dozen people about HPC at the national labs and how it is helping industry stay competitive. Messina spoke with Moniz as he toured the exhibition with a group of congressmen, national laboratory directors, and federal national laboratory representatives.

Laboratory Open House

On May 21, 2016, Argonne hosted a public open house that attracted over 10,000 people and featured campus-wide demonstrations, exhibits, and facility tours. Computing-related demos and exhibits included 3D demos on the Active Mural display in the Vis Lab and informational exhibits outside the datacenter. Four additional stations were organized around the following themes: science at the ALCF, computing hardware, Mira and the coming Theta and Aurora, and computer room safety.

Visits and Tours

The laboratory also hosts many tour groups, invited officials, and meeting attendees throughout the year. Many of these groups also request facility tours, often several times a week, which gives ALCF staff the opportunity to interact with a wide range of visitors having differing levels of expertise and varying interests.

International Delegations

On January 25, 2016, Argonne's Paul Domagala and David Martin met with Vincent Floreani, Consul General of France in Chicago; Marc Rousset, Attaché for Science & Technology; and Michel Gilbert and Jean-Pascal Saucier of Invest in France. This gathering was a follow-on meeting to educate the Consul General and French business development executives about Argonne's capabilities and to discuss how ALCF engages with industry and research partners. Brazil's National Confederation of Industry delegation visited Argonne on April 21, 2016, and met with computational scientist Roberto Paoli and Deputy Director of Operations Ti Leggett.

STEM-Related Activities

On April 1, 2016, the ALCF hosted two STEM-oriented university student organizations: Women in Science, Technology, Informatics, and Math LLC from Indiana University Bloomington, and the IEEE Chicago Section Student Branch. The IEEE-Chicago Section includes Illinois Institute of Technology, ITT Technical Institute, Concordia University, DePaul University, Northwestern University, and the University of Illinois at Chicago.

On March 18, 2016, U.S. Rep. Randy Hultgren (IL-14) visited the ALCF to present an award certificate to Oswego High School senior Jake Cirino, 14th District winner of the Congressional App Challenge in the fields of computer science and technology, for his software application, "Pollution Simulator." The competition is designed to engage student participation in STEM education fields. Hultgren serves on the Science, Space, and Technology Committee and the Subcommittees on Energy and on Research and Technology. He co-founded the Science and National Labs Caucus and also co-chairs the STEM Education Caucus.

1.3.6 Communications

Communications through Mailing Lists and Social Media

ALCF provided information to users through several electronic communication channels, including direct e-mails, custom-tailored e-mail messages via scripts, social media postings, and ALCF website postings (Table 1.6; target audiences are identified in Table 1.7). Users can opt out of the system notify and newsletter mailing lists.

Table 1.6 2016 Primary Communication Channels

Channel Name	Description	When Used/Updated
Newsbytes	HTML-formatted newsletter featuring science, facility news, recent news hits, and upcoming training events.	Monthly
Special Announcements	E-mail newsletter with information on conferences, training events, etc.—both ALCF and non-ALCF opportunities.	Ad hoc
Weekly Digest	Plain-text weekly rollup of events affecting ALCF systems and software, upcoming deadlines, and training opportunities.	Weekly
Social Media	Social media used to promote ALCF news and events.	Frequently
ALCF Website	An integrated information hub for user documentation, program and resources descriptions, user-centric events, feature stories about users, and related news.	Frequently
Custom E-mail Messages	Notification of machine status or facility availability typically formatted text-based per user and channel preference.	As needed

Table 1.7 2016 Target Audiences

Channel	Target Audience(s)
Newsbytes	Users, scientific communities, students, the general public
Special Announcements	Users, scientific communities, students, the general public
Weekly Digest	Current users on the systems with accounts
Social Media	Users, followers of ALCF, collaborators, students, scientific communities, the general public
ALCF Website	Users, collaborators, students, scientific communities, the general public
Custom E-mail Messages	Specific projects, user groups, Principal Investigators/proxies, individual users

The ALCF's monthly newsletter, Newsbytes, features science stories that highlight the outcomes of research carried out on ALCF resources or advancements made by ALCF staff and researchers in the field. This e-publication also announces training opportunities and events, allocation program announcements, and relevant news stories. Special announcements are sent out to call attention to an event or opportunity, such as the open call for participation for ATPESC.

Promotional Activities and Media Hits

ALCF published 45 science stories in 2016 on the ALCF website and in Newsbytes newsletters. All of these science stories are also channeled to Argonne's media and public relations teams. Thanks to the efforts of Argonne's media and public relations as well as ALCF direct relationships, the ALCF accrues media hits, which are stories that appear in various media outlets (print, online, TV/radio, etc.) that have passed through some form of "editorial filter," i.e., a person has made a decision to run or not run the story. In 2016, the facility posted 78 such media hits to the website. The media team uses Meltwater News public relations suite to help track media hits. This global online media monitoring company tracks articles from more than 200,000 news publications, Twitter, YouTube, Facebook, and blogs. In 2016, Meltwater captured 423 mentions of "Argonne Leadership Computing Facility" and "ALCF."

Other Publications

The ALCF produces a variety of print publications used for promotion, education, and recruitment (Table 1.8). In addition, Argonne visitors who tour the ALCF receive an informational packet tailored to their particular area(s) of interest. Most of these documents are also made available via the ALCF website. The ALCF Annual Report was recognized by the Society of Technical Communications, receiving a Distinguished Award and Best in Show from the society's Chicago Chapter, as well as an award at the society's International Summit Competition.

Table 1.8 Publications Designed for Print in 2016

Publication	Frequency	When
INCITE Poster	Yearly	January
INCITE Brochure	Yearly	November
Fact Sheet	Yearly	November
Annual Report	Yearly	March
Science Highlights	Yearly	September
Press and Visitor Packets	As Needed	As Needed
Industry Brochure	Yearly	October

Conclusion

As a user facility, ALCF is focused on ensuring the success of all facility users and customers. All of our user satisfaction ratings remain at the level expected of a leadership facility. During CY 2016, ALCF enhanced our users' onboarding experience, helped our users improve the performance of their scientific codes, made improvements to both external and internal reports, adopted best practices, continued to engage with user communities in both academia and industry, worked with users to improve application performance, improved monitoring and data analysis, automated archiving from file system to tape, prepared the facility to support our next generation platforms, volunteered in STEM-related educational efforts, invested in lab-wide and national efforts in diversity in both HPC and science communities, and enhanced overall communication efforts in various e-mail channels. As such, ALCF continues to help its users succeed by providing effective Tier 1 support, application support, resource support, outreach, and communications.

Section 2. Business Results

Is the Facility maximizing the use of its HPC systems and other resources consistent with its mission?

ALCF Response

ALCF has exceeded the metrics target for system availability, INCITE hours delivered, and capability hours delivered. For the reportable areas, such as Mean Time to Interrupt (MTTI), Mean Time to Failure (MTTF), and system utilization, ALCF is on par with the other facilities and has demonstrated exceptional performance. To assist in meeting these objectives and to improve overall operations, ALCF tracks hardware and software failures and analyzes their impact on the user jobs and metrics as a significant part of its improvement efforts.

Table 2.1 summarizes all metrics reported in this section.

Table 2.1 Summary of All Metrics Reported in the Business Results Section

Mira (Blue Gene/Q): 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	CY 2015		CY 2016	
	Target	Actual	Target	Actual
Scheduled Availability	90.0%	99.2%	90.0%	98.9%
Overall Availability	90.0%	96.3%	90.0%	94.9%
System MTTI	N/A	9.50 days	N/A	10.00 days
System MTTF	N/A	24.16 days	N/A	40.55 days
INCITE Usage	3.5B	4.0B [c]	3.5B	3.9B [e]
Total Usage	N/A	5.9B [d]	N/A	6.0B [f]
System Utilization	N/A	89.2%	N/A	90.8%
Mira INCITE Overall Capability [a]	40.0%	73.4% [c]	40.0%	76.4% [e]
Mira INCITE High Capability [b]	10.0%	31.0% [c]	10.0%	40.9% [e]

[a] Overall Capability = Jobs using ≥ 16.7 percent (8 racks, 131,072 cores) of Mira.

[b] High Capability = Jobs using ≥ 33.3 percent (16 racks, 262,144 cores) of Mira.

[c] Usage includes 9.9M core-hours from Cetus production jobs.

[d] Usage includes 30.2M core-hours from Cetus production jobs.

[e] Usage includes 3,604 core-hours from Cetus production jobs.

[f] Usage includes 24.1M core-hours from Cetus production jobs.

ALCF Resources

During CY 2016, ALCF operated one production resource, Mira. Mira is a 48K-node, 768K-core, 10 PF Blue Gene/Q with 768 TB of RAM. Mira mounts three GPFS file systems with approximately 26.5 PB of usable space and has access to the facility-wide HPSS (high-performance storage system) tape archive. Mira has an associated visualization and analysis cluster called Cooley. ALCF operated two other Blue Gene/Q systems, Cetus and Vesta.

Cetus is a 4K-node, 64K-core Blue Gene/Q with 64 TB of RAM. Cetus shares file systems with Mira. Vesta is a 2K-node, 32K-core Blue Gene/Q with 32 TB of RAM. Vesta is an independent test and development resource and shares no resources with Mira or Cetus.

In 2014, ALCF began select use of Cetus for INCITE projects with simulation runs that required nontraditional HPC workflows and has continued implementing that usage approach ever since. This deployment of Cetus allowed Mira to continue to operate as designed and enabled a new class of leadership applications to be supported.

2.1 Resource Availability

Overall availability is the percentage of time a system is available to users. Outage time reflects both scheduled and unscheduled outages. For HPC Facilities, scheduled availability is the percentage of time a designated level of resource is available to users, excluding scheduled downtime for maintenance and upgrades. To be considered a scheduled outage, the user community must be notified of the need for a maintenance event window no less than 24 hours in advance of the outage (emergency fixes). Users will be notified of regularly scheduled maintenance in advance, on a schedule that provides sufficient notification, and no less than 72 hours prior to the event, and preferably as much as seven calendar days prior. If that regularly scheduled maintenance is not needed, users will be informed of the cancellation of that maintenance event in a timely manner. Any interruption of service that does not meet the minimum notification window is categorized as an unscheduled outage. A significant event that delays a return to scheduled production will be counted as an adjacent unscheduled outage. Typically, this would be for a return to service four or more hours later than the scheduled end time. The centers have not yet agreed on a specific definition for this rare scenario.

This section reports on measures that are indicative of the stability of the system and the quality of the maintenance procedures.

2.1.1 Scheduled and Overall Availability

Mira has been in full production since April 9, 2013. In consultation with the DOE Program Manager, ALCF has agreed to metrics of 90 percent overall availability and a target of 90 percent scheduled availability (ASCR requested that all user facilities use a target of 90 percent for scheduled availability for the lifetime of the production resources). Table 2.2 summarizes the availability results.

Table 2.2 Availability Results

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	CY 2015		CY 2016	
	Target (%)	Actual (%)	Target (%)	Actual (%)
Scheduled Availability	90.0	99.2	90.0	98.9
Overall Availability	90.0	96.3	90.0	94.9

The remainder of this section covers significant availability losses, and responses to them, for both scheduled and overall availability data.

Explanation of Significant Availability Losses

This section briefly describes the causes of major losses of availability for the period January 1, 2016, through December 31, 2016, as annotated in Figure 2.1.

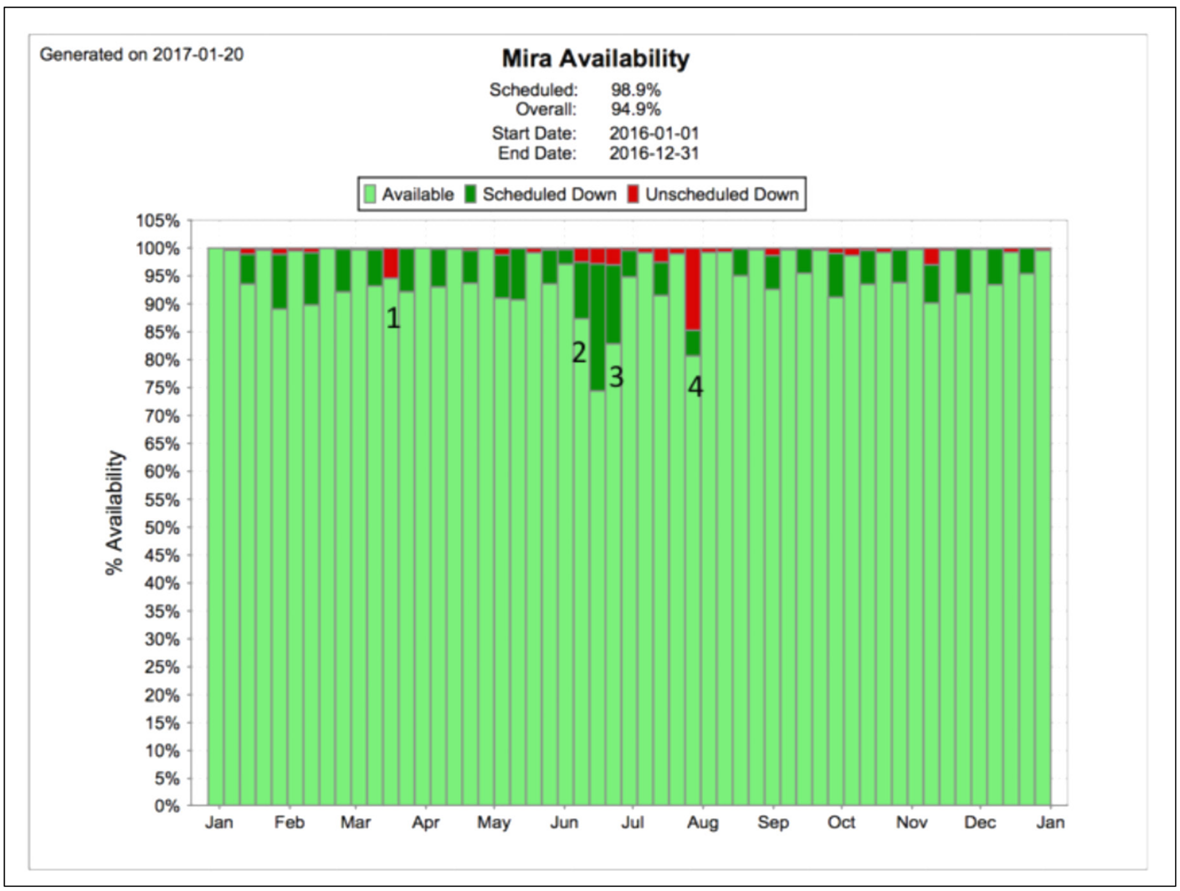

Figure 2.1 Mira Weekly Availability for CY 2016

Graph Description: Each bar in Figure 2.1 represents the average of seven days of core-hour usage. Each bar accounts for all of the time in one of three categories. The pale-green portion represents available core-hours; the darker green represents scheduled downtime for that week; and red represents unscheduled downtime. The numeric annotations are the significant losses. Each of these events is described in detail below.

Item 1: GPFS failure – March 13, 2016
On Sunday March 13, 2016, Mira-fs1 (a 7 PB scratch filesystem for Mira) hung. This event was attributable to a bug in GridScalar 3.1 (previously known as GPFS) that results in block drivers dropping I/O requests. The storage team had to restart the filesystem, and all jobs running at the time were killed. These hangs happen infrequently and are resolved by shutting down GPFS only on the server node with the "stuck/missing" I/O request. The filesystem has been monitored closely, including during off-hours, to minimize the possibility that this bug will again hang the filesystem. The filesystem will be upgraded to GridScalar 3.2.1 once all testing is complete.

Item 2: Hardware debugging after water loop interruption – June 6, 2016
On June 6, 2016, the TCS Building Data Center finalized work on pump upgrades to the Mira/Theta water loop on a scheduled maintenance day. However, there was an unexpected flow interruption late in the day that caused all of the BG/Q cabinets to SCRAM. Some hardware was damaged as a result of the SCRAMs: two nodeboards, two power modules, and a coolant monitor. The coolant monitor is notable as it took more than a week to diagnose it as faulty, and it was finally replaced two weeks after the event. This event is called out because it resulted in the need to make three downstream reservations to debug the problem with the coolant monitor. Because the coolant monitor was part of a midplane that controls half of the Mira racks clocks, the three troubleshooting reservations required half of Mira to be out of service for several hours – these are detailed next.

1st troubleshooting reservation – June 9, 2016
On June 9, 2016, bulk power enclosures (BPE) were swapped for diagnostics, and node boards were returned to service but only for a few hours before the nodes were unusable. The reservation was for 4 hours and 16 minutes on half of Mira.

2st troubleshooting reservation – June 13, 2016
On June 13, 2016, one BPE was swapped for diagnostics, and one midplane was swapped. However, again, node boards only stayed in service for a few hours during testing. The reservation was for 7 hours on half of Mira.

3rd troubleshooting reservation – June 20, 2016
On June 20, 2016, a coolant monitor was replaced in a midplane and has been working fine ever since. This reservation was for 7 hours on half of Mira. This event piggy-backed the next outage – Item 3.

Item 3: Scheduled outage for tie switch upgrades – June 17-19, 2016
On June 17–19, 2016, there was a scheduled, full power outage for the TCS Building. This outage was initiated to accomplish a reprogramming of the building bus tie switch and to work on integrating an uninterruptible power supply (UPS) for the ALCF-Lithium Theta project. The work was scheduled for the entire weekend. The ALCF Operations team brought down the ALCF resources prior to the outage early on Friday, June 17, 2016, and then brought the system back

up starting at noon on Sunday, June 19, 2016, and had it back in service by 6:45 pm. This outage lasted for 2.5 days.

Item 4: Partial site-wide outage – July 28–29, 2016
On Thursday, July 28, 2016, at 2:28 pm, the lab suffered a partial site-wide outage. The TCS Building suffered an outage when a 13.2-kV power line fell onto a 138-kV line (feeding our substation) because of high winds. Power was restored in the TCS Building at 5:35 pm. However, ALCF had already sent people home for the day because any repair work would have taken all night, and it was questionable whether the cooling and heat loads would have been stable overnight because all of the other data center tenants had left and their systems would not be generating heat. The building and all of the tenants planned to start at 6:00 am on Friday, July 29, 2016, to bring back the cooling and heat loads simultaneously. ALCF began bringing its hardware and services back at 6:45 am on July 29, 2016. It took longer than normal to return Mira to service, but that was to be expected considering how it crashed. Our team had to work on each individual cabinet to bring each back from its individual state, sometimes performing multiple, time-consuming resets. Mira was released to users at 6:33 pm. The total length of the outage was 28 hours and 5 minutes. As for damage sustained from the power outage event, a few disks were lost but there was no data loss. Mira also suffered the loss of some hardware, although the loss was limited to two pieces of hardware: one business process manager (BPM) and one data computing appliance (DCA), and these were replaced during the bring up.

2.1.2 System Mean Time to Interrupt (MTTI) and System Mean Time to Failure (MTTF)

MTTI = Time, on average, to any outage on the system, whether unscheduled or scheduled. Also known as MTBI (Mean Time Between Interrupt).
MTTF = Time, on average, to an unscheduled outage on the system.

ALCF MTTI and MTTF Summary

MTTI and MTTF are reportable values with no specific target. Table 2.3 summarizes the current MTTI and MTTF values.

Table 2.3 MTTI and MTTF Results

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	CY 2015		CY 2016	
	Target	Actual	Target	Actual
System MTTI	N/A [a]	9.50 days	N/A	10.00 days
System MTTF	N/A	24.16 days	N/A	40.55 days

[a] N/A = Not applicable.

Mira currently functions on a biweekly maintenance schedule. ALCF takes the machine out of service every other Monday to perform Blue Gene driver upgrades, hardware replacements, operating system (OS) upgrades, etc. Further, while Mira is out of service, ALCF uses that opportunity to perform other potentially disruptive maintenance such as facilities power and cooling work and storage systems upgrades and patching. ALCF's biweekly maintenance schedule caps MTTI at 14 days, but does not directly affect MTTF.

2.2 Resource Utilization

The following sections discuss system allocation and usage, total system utilization percentage, and capability usage. For clarity, *usage* is defined as resources consumed in units of core-hours. *Utilization* is the percentage of the available core-hours used (i.e., a measure of how busy the system was kept when it was available).

2.2.1 Total System Utilization

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

Utilization is a reportable value with no specific target. A rate of 80 percent or higher is generally considered acceptable for a leadership-class system. Table 2.4 summarizes ALCF utilization results, and Figure 2.2 shows Mira system utilization over time by program.

Table 2.4 System Utilization Results

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	CY 2015		CY 2016	
	Target	Actual	Target	Actual
System Utilization	N/A [a]	89.2%	N/A	90.8%

[a] N/A = Not applicable.

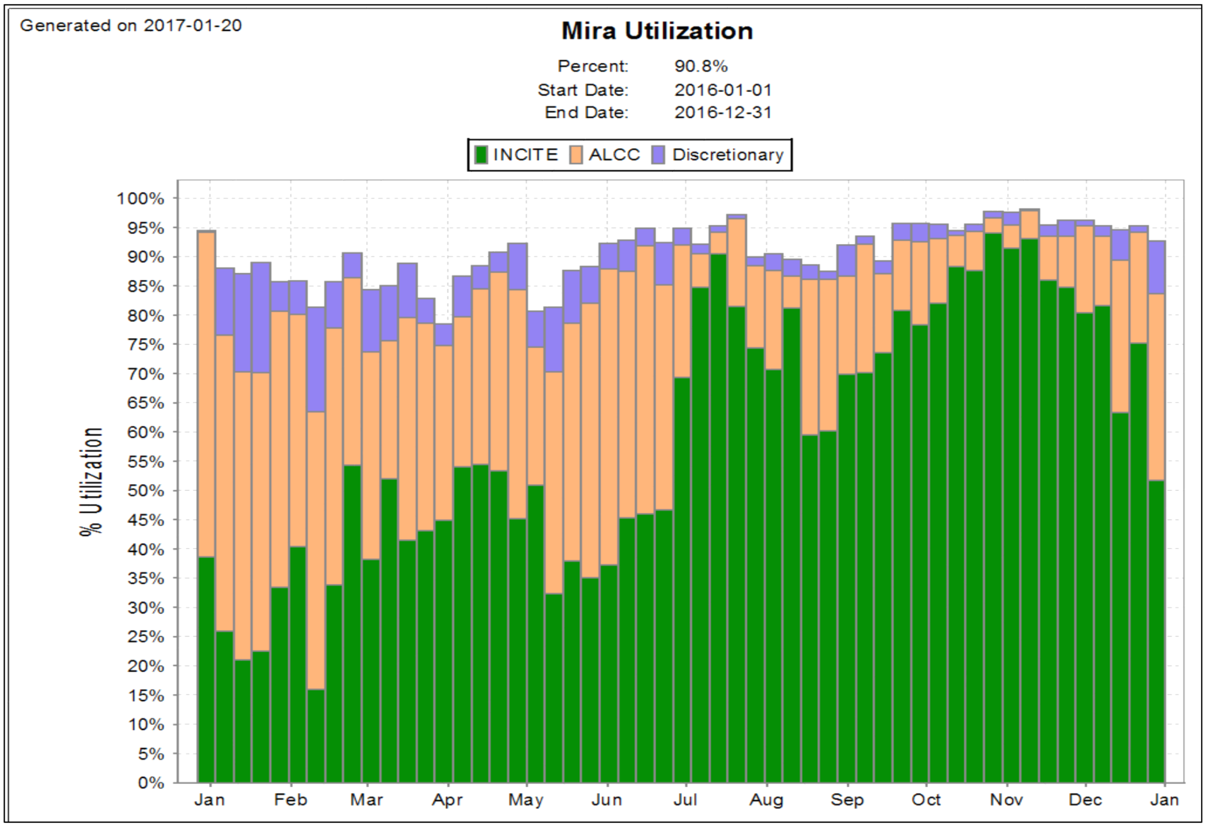

Figure 2.2 System Utilization over Time by Program

The system utilization for Mira was 90.8 percent for its 2016 production period of January 1, 2016, through December 31, 2016.

Table 2.5 shows how Mira's system hours were allocated and used by allocation source. Multiplying the theoretical hours by availability and utilization values that were agreed upon with ALCF's DOE Program Manager determines the hours available. Of the hours available, 60 percent is allocated to the INCITE program, up to 30 percent is available for the ALCC program allocations, and 10 percent is available for Director's Discretionary (DD) allocations. The ALCC program runs from July through June, so to arrive at allocated values for the calendar year, half of the hours are assigned to each year. The allocated values for the DD allocations appear higher than expected because they represent a rolling allocation. A majority of DD projects are exploratory investigations, so the time allocations are often not used in full. DD allocations are discussed in detail in Section 3.3. In CY 2016, ALCF delivered a total of 6.0 billion core-hours on Mira.

Table 2.5 Core-Hours Allocated and Used by Program

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM						
	CY 2015			CY 2016		
	Allocated	Used		Allocated	Used	
	Core-hours	Core-hours	%	Core-hours	Core-hours	%
INCITE	3.6B	4.0B [a]	66.7%	3.5B	3.9B [d]	65.5%
ALCC	1.7B	1.6B [b]	27.0%	1.7B	1.7B [e]	28.4%
DD	858.3M	373.2M [c]	6.3%	1.0B	365.5M [f]	6.1%

[a] Usage includes 9.9M core hours from Cetus production jobs.

[b] Usage includes 10.4M core hours from Cetus production jobs.

[c] Usage includes 9.9M core hours from Cetus production jobs.

[d] Usage includes 3,604 core hours from Cetus production jobs.

[e] Usage includes 2.9M core hours from Cetus production jobs.

[f] Usage includes 21.2M core hours from Cetus production jobs.

Summary: For CY 2016, the system usage and utilization values were in line with general expectations.

2.3 Capability Utilization

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

On Mira, capability is defined as using greater than 16.7 percent of the machine. Historically, capability has been defined as using greater than 20 percent of the machine. However, 20 percent of Mira would be 9.6 racks, which is not a viable configuration. Hence, the Mira capability metric is split into two parts. Overall Capability requires that a minimum of 40 percent of the INCITE core-hours be run on eight racks or more (16.7 percent), and High Capability requires a minimum of 10 percent of the INCITE core-hours be run on 16 racks or more (33.3 percent). Table 2.6 and Figure 2.3 show that ALCF has substantially exceeded these metrics set for INCITE. Although no targets are set, data are also provided in the table for ALCC and DD projects as reference, and Figure 2.4 shows the overall distribution of job sizes over time.

Table 2.6 Capability Results

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM						
	CY 2015			CY 2016		
Capability Usage	Total Hours	Capability Hours	Percent Capability	Total Hours	Capability Hours	Percent Capability
INCITE Overall	4.0B	2.9B	73.4%	3.9B	3.0B	76.4%
INCITE High	4.0B	1.2B	31.0%	3.9B	1.6B	40.9%
ALCC Overall	1.6B	809.9M	50.5%	1.7B	1.0B	60.0%
ALCC High	1.6B	298.5M	18.6%	1.7B	327.0M	19.2%
Director's Discretionary Overall	373.2M	202.4M	54.2%	365.5M	158.1M	43.3%
Director's Discretionary High	373.2M	127.2M	34.1%	365.5M	78.7M	21.5%
TOTAL Overall	5.9B	3.9B	66.0%	6.0B	4.2B	69.7%
TOTAL High	5.9B	1.6B	27.8%	6.0B	2.0B	33.6%

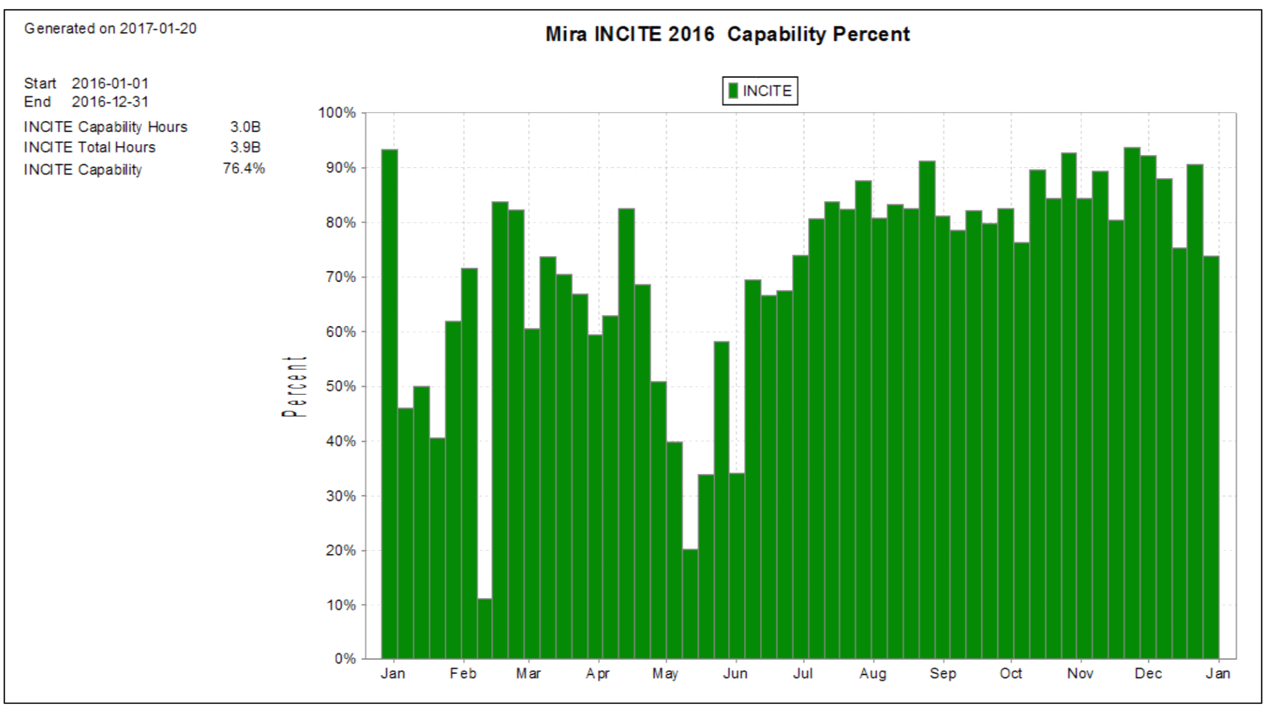

Figure 2.3 Mira INCITE Overall Capability

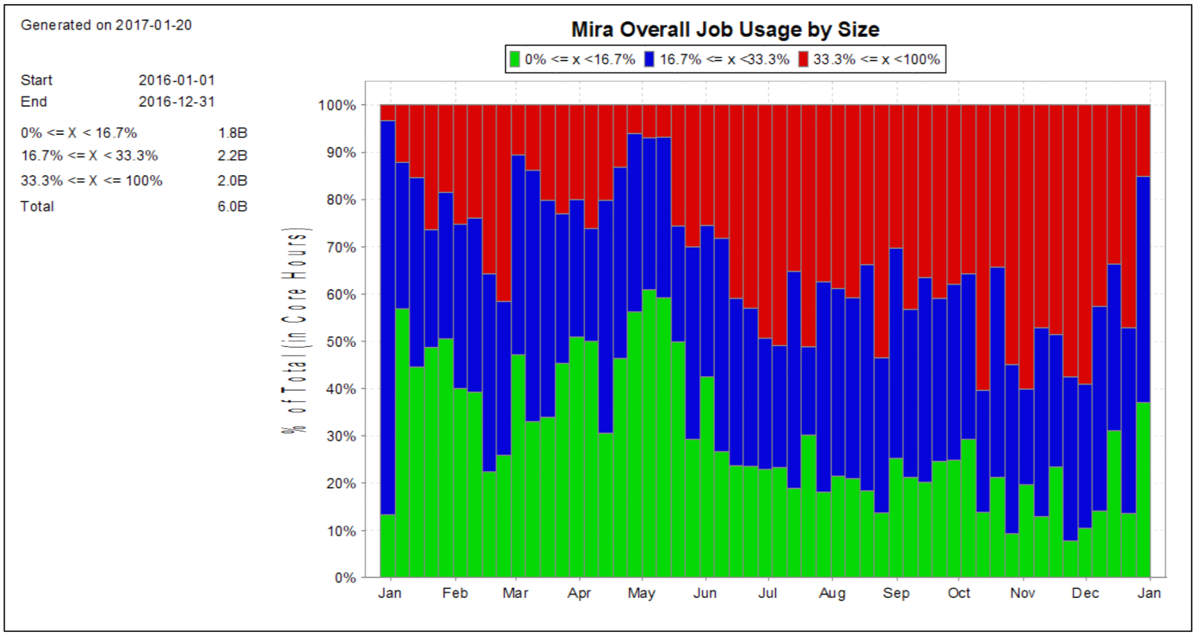

Figure 2.4 Mira Job Usage by Size

2.4 Financial and Human Resources

Financial data and staffing distribution are covered in detail in the onsite review. A summary is provided here. ALCF carefully plans the use of its available resources, maximizing the impact of its staffing distribution through strategic use of matrixed employees and contractors where consistent with its mission. Annual and multi-year budgets are reviewed with ASCR and the Argonne Site Office (DOE-ASO) on a regularly scheduled basis. Actual costs are tracked closely against approved budgets. Ongoing communication with ASCR and DOE-ASO includes a weekly conference call and a monthly staffing update. This line of communication allows for a discussion of the best use of ALCF resources. In FY 2016, equipment purchases and the Theta lease down payment were coordinated efforts with the ASCR program office.

Conclusion

The ALCF is maximizing the use of its HPC systems and other resources consistent with its mission. The facility has exceeded the metrics of system availability, INCITE hours delivered, and capability hours delivered. For the reportable areas—MTTI, MTTF, and utilization—ALCF is on par with OLCF and NERSC, and the values reported are reasonable. These measures are summarized in Table 2.7.

The ALCF closely tracks hardware and software failures and their impact on user jobs and metrics. These data are used as a significant factor in the selection of troubleshooting efforts and improvement projects. In CY 2016, this regular failure analysis has continued to drive code changes to Cobalt, ALCF's job scheduler, and has provided details to support the debugging of storage systems, as noted in Section 2.1.1 under "Explanation of Significant Availability Losses."

Table 2.7 Summary of All Metrics Reported in the Business Results Section

Mira (Blue Gene/Q): 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	CY 2015		CY 2016	
	Target	Actual	Target	Actual
Scheduled Availability	90.0%	99.2%	90.0%	98.9%
Overall Availability	90.0%	96.3%	90.0%	94.9%
System MTTI	N/A	9.50 days	N/A	10.00 days
System MTTF	N/A	24.16 days	N/A	40.55 days
INCITE Usage	3.5B	4.0B [c]	3.5B	3.9B [e]
Total Usage	N/A	5.9B [d]	N/A	6.0B [f]
System Utilization	N/A	89.2%	N/A	90.8%
Mira INCITE Overall Capability [a]	40.0%	73.4% [c]	40.0%	76.4% [e]
Mira INCITE High Capability [b]	10.0%	31.0% [c]	10.0%	40.9% [e]

[a] Overall Capability = Jobs using ≥ 16.7 percent (8 racks, 131,072 cores) of Mira.

[b] High Capability = Jobs using ≥ 33.3 percent (16 racks, 262,144 cores) of Mira.

[c] Usage includes 9.9M core-hours from Cetus production jobs.

[d] Usage includes 30.2M core-hours from Cetus production jobs.

[e] Usage includes 3,604 core-hours from Cetus production jobs.

[f] Usage includes 24.1M core-hours from Cetus production jobs.

This page intentionally left blank.

Section 3. Strategic Results

Are the processes for science engagement and outreach effective and do these processes enable scientific achievements consistent with the Department of Energy strategic goals?

ALCF Response

The science accomplishments of INCITE, ALCC, and DD projects clearly demonstrate ALCF's impact in supporting scientific breakthroughs. ALCF staff members have worked effectively with individual project teams to adapt their simulation codes to run efficiently in a high-performance computing environment and have enabled scientific achievements that would not have been possible otherwise.

In this section, ALCF reports:

- Science Output (publications);
- Scientific Accomplishments (true accomplishments, not just milestones)
- Allocation of Facility Director's Reserve Computer Time; and
- Joint Laboratory for System Evaluation (JLSE)

3.1 Science Output

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

Table 3.1 shows the breakdown by journal of refereed publications based (at least in part) on the use of ALCF resources. The Nature Journals category refers to two publications in Nature and three in Nature Communications. In addition to these journals, ALCF users published in journals such as *Science*, *Proceedings of the National Academy of Sciences (PNAS)*, *Physical Review Letters, and The International Conference for High Performance Computing, Networking and Analysis (SC'16)*.

Table 3.1 Summary of Refereed Publications

Nature Journals	Science	PNAS	Physical Review Letters	SC '16	Total 2016 Publications
5	1	1	11	4	254

3.2 Scientific Accomplishments

The Facility highlights a modest number (top five) of significant scientific accomplishments of its users, including descriptions for each project's objective, the implications of the results achieved, the accomplishment itself, and the facility's actions or contributions that led to the *accomplishment. The accomplishment slides (which are NOT limited to five) should include the allocation, amount used, and a small bar graph indicating size of jobs. (See ALCF 2011 slides as example.)*

LCFs should include tables/charts comparing time allocated to time used by projects. Because of their much larger volume of projects, NERSC should include a chart summarized by SC program.

Each science highlight includes a box with a bar graph. The top line indicates the machine used, program (INCITE, ALCC, DD), and the year(s) of the allocation(s). The second line lists the total core-hours allocated to the project and, in parentheses, the core-hours used. The graph shows the core-hour breakdown for each project, by the percentage of the machine used. The breakdown is based on the ALCF capability metric detailed in Section 2. The bar on the left (where present) represents simulation runs below capability, the middle bar (where present) represents runs at the first capability threshold up to but not including the second threshold, and the bar on the right (where present) represents runs at the highest capability.

3.2.1 Reactive MD Simulations of Electrochemical Oxide Interfaces at Mesoscale (INCITE 2015)

Subramanian K. R. S. Sankaranarayanan, Argonne National Laboratory

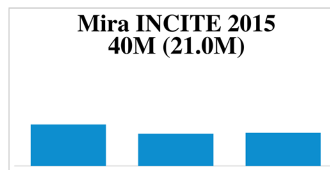

The discovery of this ultra-durable, self-lubricating tribofilm (a film that forms between moving surfaces) can have profound implications for the efficiency and durability of future engines and other moving metal parts that can be made to develop their own self-healing, diamond-like carbon (DLC) tribofilms.

The original discovery occurred when scientists decided to see what would happen when a small steel ring was coated with a catalytically active nanocoating (tiny molecules of metals that promote chemical reactions to break down other materials) and then subjected to high pressure and heat using a base oil without the complex additives of modern lubricants. When they examined the ring after the endurance test, they did not see the expected rust and surface damage, but an intact ring with an odd blackish deposit on the contact area. Viewing the deposit using high-powered optical and laser Raman microscopes, the experimentalists realized the deposit was a tribofilm of diamond-like carbon, similar to several other DLCs developed at Argonne in the past. Further experiments, led by postdoctoral researcher Giovanni Ramirez, revealed multiple catalytic coatings that yield DLC tribofilms. Experiments showed that the coatings interact with the oil molecules to create the DLC film, which adheres to the metal surfaces. The scientists realized the film was developing spontaneously between the sliding surfaces and was replenishing itself: next, they needed to understand why and how.

To provide the theoretical understanding of what the tribology team was discovering in its experiments, they turned to Argonne's theoretical scientists, who used the immense computing power of ALCF's 10-petaflop supercomputer, Mira. They ran large-scale reactive MD simulations using the LAMMPS code to understand what was happening at the atomic level, and determined that the catalyst metals in the nanocomposite coating were stripping hydrogen atoms from the hydrocarbon chains of the lubricating oil, then breaking the chains down into smaller segments. The smaller chains joined together under pressure to create the highly durable DLC tribofilm (Figure 3.1).

Mira simulations also allowed researchers to look beyond the current study by performing virtual tests of other potential catalysts (other metals and hydrocarbons in coatings and oils) for their "self-healing" properties in a high-temperature, high-pressure engine environment.

IMPACT: There are enormous implications for the new tribofilm that can potentially increase the efficiency and reliability of engines. Manufacturers already use many different types of coatings – some developed at Argonne – for metal parts in engines and other applications. The problem is that those coatings are expensive and difficult to apply, and once they are in use, they only last until the coating wears through. The new catalyst allows the tribofilm to be renewed continually during operation. In addition, it could allow manufacturers to reduce, or possibly eliminate, some of the modern anti-friction and anti-wear additives in oil. These additives can decrease the efficiency of vehicle catalytic converters and can be harmful to the environment because of their heavy metal content.

ALCF Contributions: With the help of ALCF staff in 2015, a team of domain and computational scientists worked to improve LAMMPS performance. The improvements targeted several parts of the code, including the ReaxFF module, an add-on package used to model the chemical reactions occurring in the system.

In collaboration with researchers from IBM, Lawrence Berkeley National Laboratory (LBNL), and Sandia National Laboratories, ALCF optimized LAMMPS by replacing the MPI point-to-point communication with MPI collectives in key algorithms, making use of MPI I/O, and adding OpenMP threading to the ReaxFF module. These enhancements doubled the code's performance.

Contributors to the code optimization work included Paul Coffman, Wei Jiang, Chris Knight, and Nichols A. Romero from the ALCF; Hasan Metin Aktulga from LBNL (now at Michigan State University); and Tzu-Ray Shan from Sandia (now at Materials Design, Inc.).

Publication: Erdemir, A., G. Ramirez, O. L. Eryilmaz, B. Narayanan, Y. Liao, G. Kamath, and S. K. R. S. Sankaranarayanan. "Carbon-based Tribofilms from Lubricating Oils." *Nature,* 2016, vol. 536, pp. 67–71. **DOI**: 10.1038/nature18948.

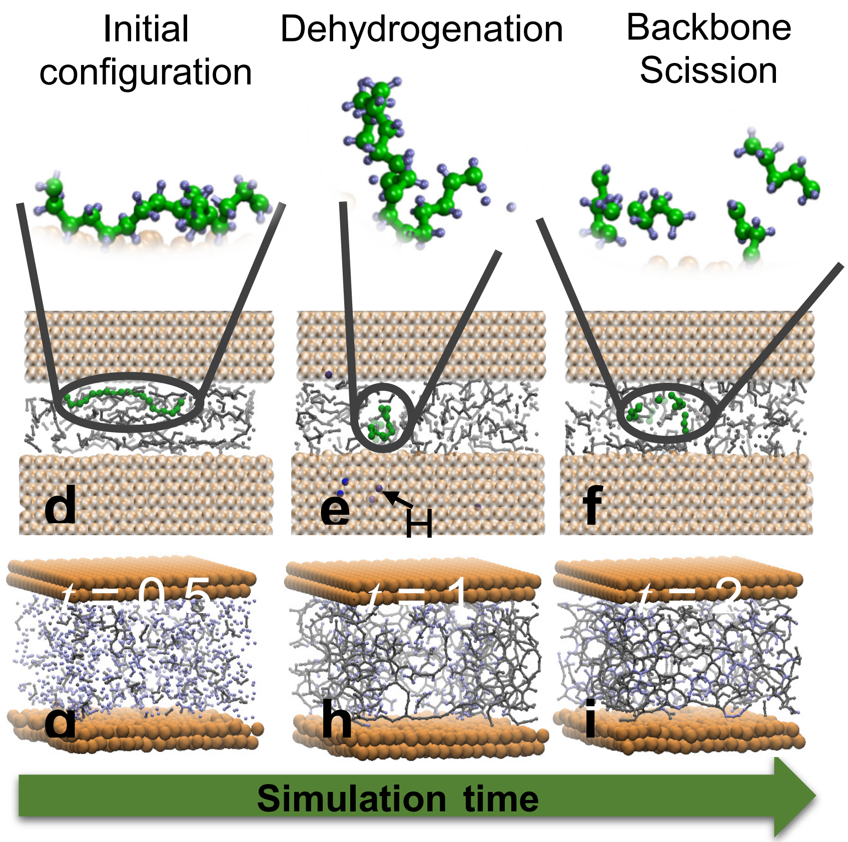

Figure 3.1 Reactive molecular dynamics (RMD) simulations show that the catalytic action of Cu (light orange) causes dehydrogenation and scission of long linear olefin chains (green and purple) into shorter chain hydrocarbons (d–f). These dehydrogenated short-chain hydrocarbons recombine to nucleate and grow into a DLC tribofilm (g–i). The DLC tribofilm formed on-demand reduces the friction drastically. The reaction pathways predicted by RMD simulations are confirmed by density functional theory (DFT) calculations. Image credit: modified figure and caption from *Nature* article.

3.2.2 Frontiers in Planetary and Stellar Magnetism through High-Performance Computing (INCITE 2015–2016)

Jonathan Aurnou, University of California, Los Angeles

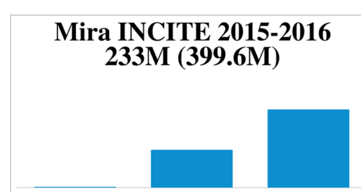

Most stars and planets harbor the propensity for generating magnetic fields through a mechanism known as dynamo action. Solar and terrestrial magnetism plays a particularly important role in our modern technological society, such as in space-borne communications, space flight, and geographical exploration. Understanding and predicting the potential long- and short-term impacts of these phenomena on the Earth requires a high-fidelity picture of the dynamo processes. Magnetic fields of astrophysical bodies also provide important clues concerning the nature of their interior dynamics, as well as their nearby space environments. In Jupiter, for instance, magnetic field observations coupled with computer simulations may provide insight into atmospheric structures and assist imaging by the National Aeronautics and Space Administration's (NASA's) spacecraft Juno.

The principle challenges facing computational dynamo studies are resolution and time integration, which demand scalable software and massively parallel computing. Solar scientists from the University of California, Los Angeles; University of California, Davis; and University of Colorado, Boulder implemented state-of-the-art algorithms for dynamo simulations. With an INCITE allocation on Mira, these scientists endeavored to refine their simulation models and explore the energy mechanisms of solar convection cells and atmospheric features of giant planets through Rayleigh, a highly scalable code.

Large-scale magnetohydrodynamics (MHD) simulations that combine two advanced algorithms, the parallelized spectral method and polynomial projection, have been carried out on Mira. The parallelized spectral method is designed to minimize process communications and increase the speed and accuracy of convergence, whereas the polynomial projection algorithm seeks to increase the speed and accuracy of spherical function operations.

The work of this INCITE project has revealed a mechanism of large-scale convection power suppression in the sun and has estimated its convective velocity. This finding provides a picture of how deep convection might give rise to the largest distinct spatial scale of convection observed in the solar photosphere at an unprecedented level of detail. The researchers have also developed high-resolution simulations of the atmospheric structure of giant planets.

IMPACT: Results from this project reveal the mechanism of large convection cell formation on the solar photosphere, lending insight into the origin of cycles of sunspot activities that affect satellite-based communications. The simulated atmospheric structure of Saturn is similar to those imaged by the Cassini spacecraft (Figure 3.2), and high-resolution simulation clarifies features of Jupiter's poles that NASA's spacecraft Juno has not been able to image.

ALCF Contributions: ALCF staff member Wei Jiang worked with the Rayleigh code developers to redesign the initialization phase to drastically reduce the number of MPI point-to-point calls and make use of MPI collectives. For capability-scale jobs at 16 or 32 racks, the initialization phase was sped up by 5 to 10 times. Wei also attended the Computational Infrastructure for Geodynamics (CIG) workshop in 2016 to provide on-site technical support for new and potential Rayleigh users.

Publications: (1) Featherstone, N. A., and B. W. Hindman. "The Spectral Amplitude of Stellar Convection and Its Scaling in the High-Rayleigh-number Regime." *Astrophys. J.*, February 2016, vol. 818, no. 1, article 32. **DOI**: 10.3847/0004-637X/818/1/32; (2) Matsui, H., et al. "Performance Benchmarks for a Next-Generation Numerical Dynamo Model." *Geochem. Geophys. Geosys*., May 2016, vol. 17, no. 5, pp. 1586–1607. **DOI**: 10.1002/2015GC006159; and (3) Featherstone, N. A., and B. W. Hindman. *Astrophys. J. Letters*, in press.

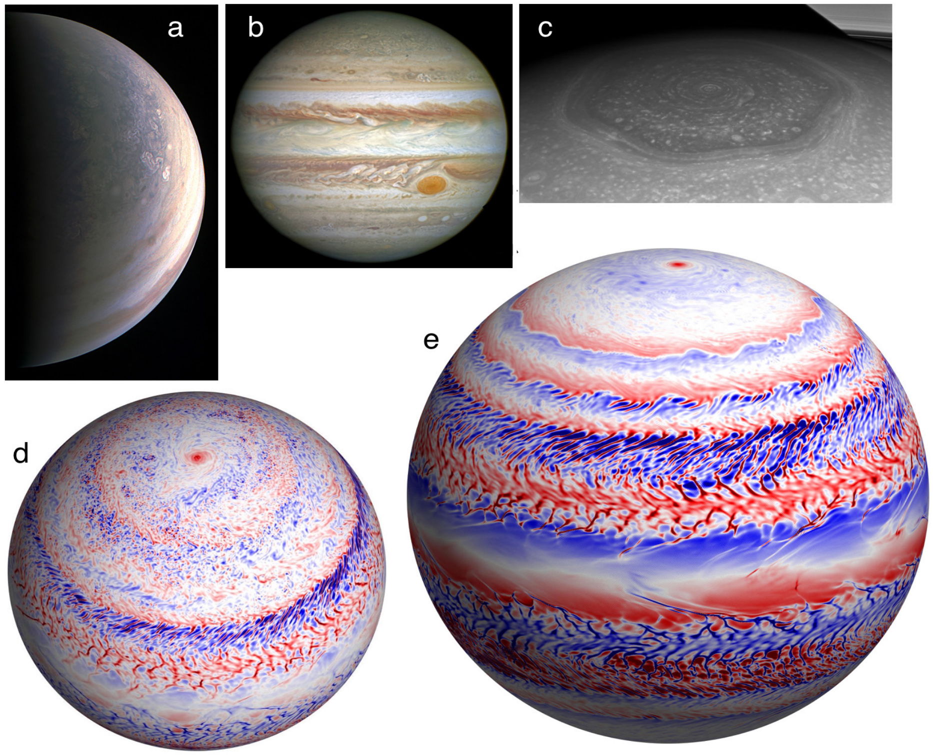

Figure 3.2 Panel (a) polar view of Jupiter from the August 27, 2016, orbital pass of Juno; (b) Jupiter in opposition from Hubble Space Telescope; (c) Cassini image of Saturn's North pole showing the hexagon and polar cyclonic vortex; (d) and (e), images of the radial vorticity (red is cyclonic and blue is anticyclonic in the northern hemisphere) near the outer boundary of the simulations.

3.2.3 Toward Breakthroughs in Protein Structure Calculation and Design (INCITE 2016)

David Baker, University of Washington

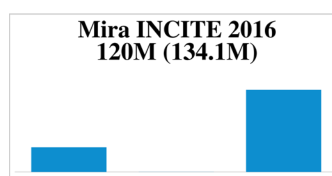

Proteins are large, complex molecules that drive virtually all cellular functions in living organisms. With the emergence of computational protein design tools, researchers have acquired the ability to design proteins with targeted applications, such as binding to other proteins involved in disease or catalyzing industrially useful chemical reactions. A research team from the University of Washington is using Mira to develop and apply new computational methods aimed at enhancing state-of-the-art protein structure prediction and design capabilities.

Most computational protein design to date has focused on proteins that can be genetically encoded and produced using transgenic bacteria or yeast cells. However, another largely unexplored realm exists: chemically synthesized mini-proteins (or *peptides*) that can be built from natural or nonnatural building blocks, paving the way to many new structural and functional possibilities that proteins made strictly from the natural amino acid building-blocks made by living cells cannot acquire. Small, synthetic peptides could potentially combine the advantageous properties of existing small-molecule medications — namely, ease of administration, the ability to cross biological barriers like the gut-blood or blood-brain barriers, evasion of the human immune system, and resistance to protein-degrading enzymes in the human body — with the advantages of large-protein therapeutics (which can bind to their targets with very high affinity and specificity, resulting in a desired therapeutic effect with few side effects).

As is the case for protein design, the key to designing a peptide drug is creating a molecule with a rigid structure that perfectly complements the shape and charge of a binding pocket in a target protein (e.g., a pathogen enzyme to be inhibited). However, peptide design presents more challenges than protein design for several reasons. Smaller size requires greater precision; simultaneously, the smaller number of atoms in a peptide produce fewer opportunities for researchers to engineer stabilizing interactions that will hold a peptide in a unique, rigid conformation to complement its target. Although chemical cross-links can help to address this problem, they also increase the difficulty of another problem — that of sampling accessible conformations of a peptide. This problem is already made difficult by the inclusion of nonnatural peptide building blocks, which prohibit natural proteins from being used as templates to guide conformational sampling. To address these challenges, the researchers have developed powerful new algorithms that have opened up this exciting new field.

The supercomputing power of Mira has allowed the researchers to sample enormous numbers of cross-linked peptide conformations, both during design and in a post-design validation phase. This high sampling level has enabled the researchers to create the first computationally designed peptides with experimentally verified, rigid structures – and has opened the door to a new frontier in drug development.

IMPACT: This project is advancing peptide structure modeling capabilities and has enabled the design of novel peptides with structures not found in nature (Figure 3.1). This development, in turn, is enabling the design of therapeutic peptides with shapes that perfectly complement target molecules for diseases such as Ebola, human immunodeficiency virus (HIV), antibiotic-resistant bacterial infection, and Alzheimer's. Nonnatural peptides represent a new class of treatments that have the potential for achieving greater efficacy while causing fewer side effects than are associated with traditional medicines. In addition, the computational tools can be used to design peptide catalysts and enzymes for various environmental, energy, and industrial applications.

ALCF Contributions: ALCF staff member Yuri Alexeev helped port the Rosetta code to Blue Gene/Q. As part of this effort, the code was benchmarked and compiled with various compilers, including the IBM C++ compiler. Yuri collaborated with the Baker group to improve its scalability on Mira. Rosetta now supports the use of multiple servers (masters) instead of a single server (master) for creation and delegation of tasks. The multiple server algorithm will allow Rosetta to, in principle, scale to the whole of Mira. So far, their production runs have used up to 24 racks of Blue Gene/Q. In collaboration with IBM software engineer Vipin Sachdeva, a part of the code was threaded for the Blue Gene/Q architecture. This ongoing effort will also help the team prepare Rosetta for Theta, the new Intel KNL supercomputer at ALCF.

Publication: Bhardwaj, G., V. K. Mulligan, C. D. Bahl, J. M. Gilmore, P. J. Harvey, O. Cheneval, G. W. Buchko, S. V. S. R. K. Pulavarti, Q. Kaas, A. Eletsky, P.-S. Huang, W. A. Johnsen, P. Greisen, G. J. Rocklin, Y. Song, T. W. Linsky, A. Watkins, S. A. Rettie, X. Xu, L. P. Carter, R. Bonneau, J. M. Olson, E. Coutsias, C. E. Correnti, T. Szyperski, D. J. Craik, and D. Baker. "Accurate De Novo Design of Hyperstable Constrained Peptides." *Nature,* 2016, vol. 538, no. 7625, pp. 329–335 (20 October 2016). **DOI:** 10.1038/nature19791.

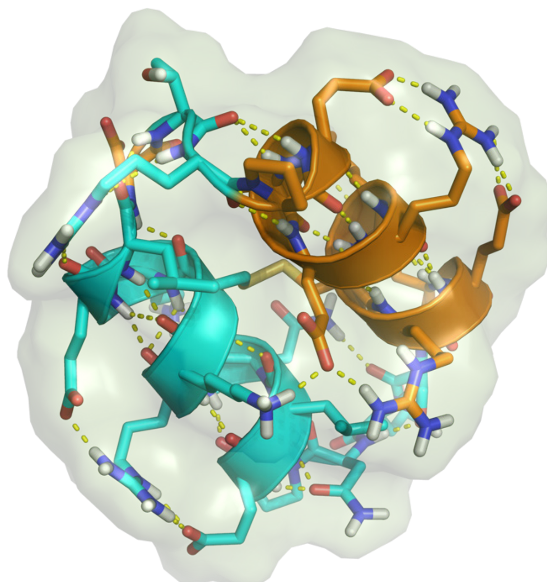

Figure 3.3 A synthetic peptide made from a mixture of natural (left-handed) amino acids, shown in cyan, and their nonnatural mirror-image counterparts (right-handed amino acids), shown in orange. The folded structure of this peptide features helices of opposite handedness packing against one another in a fold never observed in nature. Peptides of this sort have enormous potential as medications that will combine the advantages of traditional small-molecule drugs and large-protein therapeutics.

3.2.4 Anomalous Density Properties and Ion Solvation in Liquid Water: A Path-Integral *Ab Initio* Study (ALCC 2015–2016)

Robert A. DiStasio, Jr., Cornell University

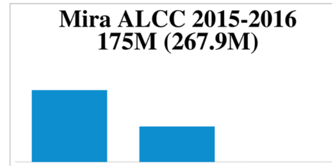

With the advent of precise laboratory techniques, it is now possible to achieve greater resolution of observables at atomic scales; thereby enabling the investigation of one of the most intriguing phenomena: the binding of polarizable structures at long-range distances (between 5 to 20 nanometers), also known as dispersion or van der Waals (vdW) forces. In the microscopic world, dispersion interactions are categorized as "weak"; nevertheless, they are the driving force for molecular assembly, responsible for the mechanical properties of polymers, the stability of proteins, and the adsorption of gases on surfaces — they even allow geckos to climb glassy surfaces (or hang on ceilings), among many other cases.

Theoretical chemists and physicists have long been challenged to find better models to explain these weak attractive forces. In the last decade, with hundreds of articles published on this topic, new models have emerged to describe dispersion interactions for quantum mechanical simulations. Some of the most successful models are based on fragment approximations and atomic pairwise interactions. Those approaches have been suitable to measure the long-range interactions of small- and medium-size molecules. However, these approaches failed to accurately describe the attractions among large structures, which result in nonphysical behavior owing to the accumulation of errors. Dr. Robert DiStasio, Jr.'s team proposed a first-principles wavelike model to quantify the dispersion interactions by taking into account the collective nonlocal, many-body contributions that give rise to the instantaneous fluctuations of charge densities. This wavelike model could come to be considered a landmark achievement in driving how scientists conceptualize the rules that govern interactions at the nanoscale.

With the help of ALCF supercomputers, this team applied and tested its wavelike model on a diverse set of nanostructures with many shapes and sizes, some of them composed of many thousands of atoms. For instance, team members analyzed the interaction energies among carbon nanotubes (Figure 3.4) and studied the attractions between nanowires and biomolecules. This model was also tested to evaluate the interactions of molecules at different finite distances. The large number of experiments performed during this project provided a strong support and validation to the model. Furthermore, these calculations helped to observe for the first time an enhancement of attractive forces at large separations for extended systems.

IMPACT: This project developed a new wavelike model based on collective many-body contributions to understand long-distance weak dispersive forces. Moving forward, this work will change how theorists include vdW interactions in their simulations.

ALCF Contributions: ALCF staff member Álvaro Vázquez-Mayagoitia assisted in porting and optimizing the codes used in this project; in particular, the Car-Parinello code (CPV) of the Quantum-Espresso package is 40% faster than the original code. Adam Scovel and Paul Rich assisted this project by packaging its jobs into larger ensemble runs.

Publication: Ambrosetti, A., N. Ferri, R. A. DiStasio Jr., and A. Tkatchenko. "Wavelike Charge Density Fluctuations and van der Waals Interactions at the Nanoscale." *Science*, 2016, vol. 351, pp. 1171–1176.

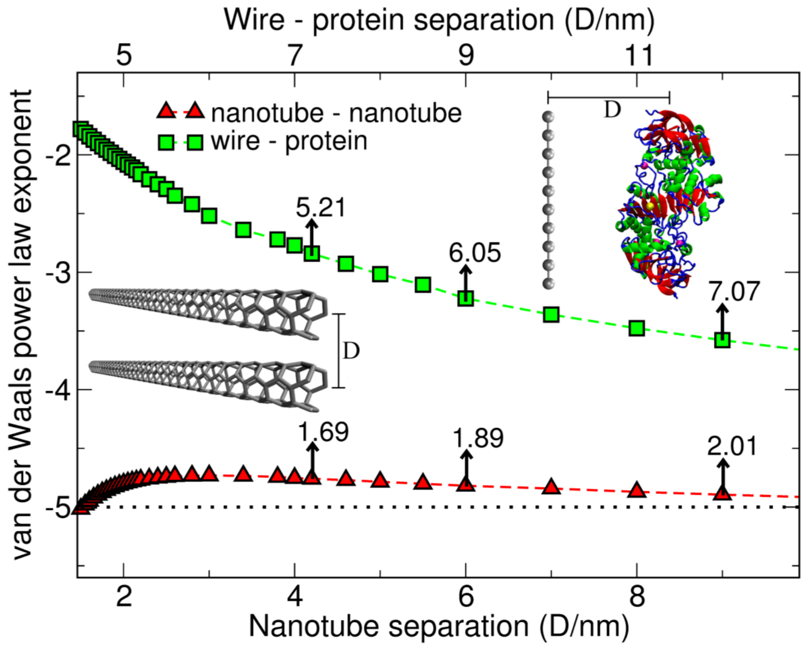

Figure 3.4 Interaction energy power law exponents computed using the many-body dispersion (MBD) model for two parallel (3, 3) carbon nanotubes and a wire–protein nanostructure as a function of the respective center-of-mass separations D (in nm).

3.2.5 Computing 3D Structures of RNA from Small-angle X-ray Scattering Data and Secondary Structures (DD 2015)

Yun-Xing Wang, National Cancer Institute

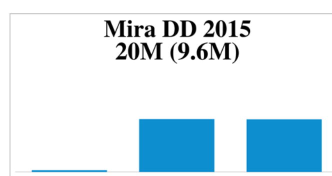

Ribonucleic acid (RNA) plays a critical role in regulating cellular processes, making it an important area of research for cancer studies. The key to understanding RNA depends on knowledge of its 3D structures, but such structures are difficult to ascertain with conventional methods. Researchers from the National Cancer Institute are using ALCF computing resources to develop a novel approach for determining RNA structures that would greatly improve our understanding of RNA functions for cancer treatment.

The research team developed a robust algorithm and computational program called RS3D to calculate 3D structures of RNA using small-angle X-ray scattering (SAXS) data and known secondary structures as input. Starting from a glob model at a nucleotide level of resolution, the algorithm carries out natural hierarchical moves based on the structural composition of

RNA. Each move is guided toward improving the SAXS data fit and long-range interactions, if available. To conduct these types of calculations, the researchers must generate tens of thousands of structures for each type of RNA. Computing the RNA structures with various folds and complexities necessitates the use of a petascale supercomputer like Mira. In collaboration with ALCF researchers, the team has scaled and optimized RS3D to improve its performance on Mira.

The team published a paper in the journal *Methods* detailing their computational approach to model RNA structures using SAXS data. They also validated the RS3D method extensively using 17 different RNA structures that represent a wide variety of folding architectures available in the current structural database (Figure 3.5). By computing 3D topological RNA structures with diverse types of junctions and structural complexities, the team demonstrated the utility and robustness of RS3D, using simulated as well as experimental SAXS data. The method's novel features include conceptual simplicity that incorporates secondary structure motifs and the flexibility to integrate a variety of tertiary interactions that users can obtain from biochemical and biophysical data.

This project is developing a high-throughput method that can determine the structure of RNA in a diverse folding landscape to fill in the gap between known RNA sequences and their explicit 3D structures. This work has significant implications for understanding the structural basis of RNA biology, and thus in the advancement of RNA therapeutics.

IMPACT: This method is a significant improvement over existing methods that seek to determine the 3D structure of RNA. As proof of concept, this method was used to determine the top 10 candidate structures for 17 different RNA species. The results from this study may have a significant impact to our knowledge about RNA structural biology, leading to greater understanding of RNA functions and activities in shutting down or modulating gene expression of cancer diseases.

ALCF Contributions: ALCF staff member Wei Jiang parallelized the RS3D code with MPI, optimized the existing OpenMP implementation, and rewrote key functions using QPX intrinsics. These enhancements led to a 4x speed-up of the RS3D code.

Publications: Bhandari, Y. R., W. Jiang, E. A. Stahlberg, J. R. Stagno, and Y. X. Wang. "Modeling RNA Topological Structures using Small Angle X-ray Scattering." *Methods*, July 1, 2016, vol. 103, pp. 18–24. Epub: June 2, 2016. **DOI:** 10.1016/j.ymeth.2016.04.015.

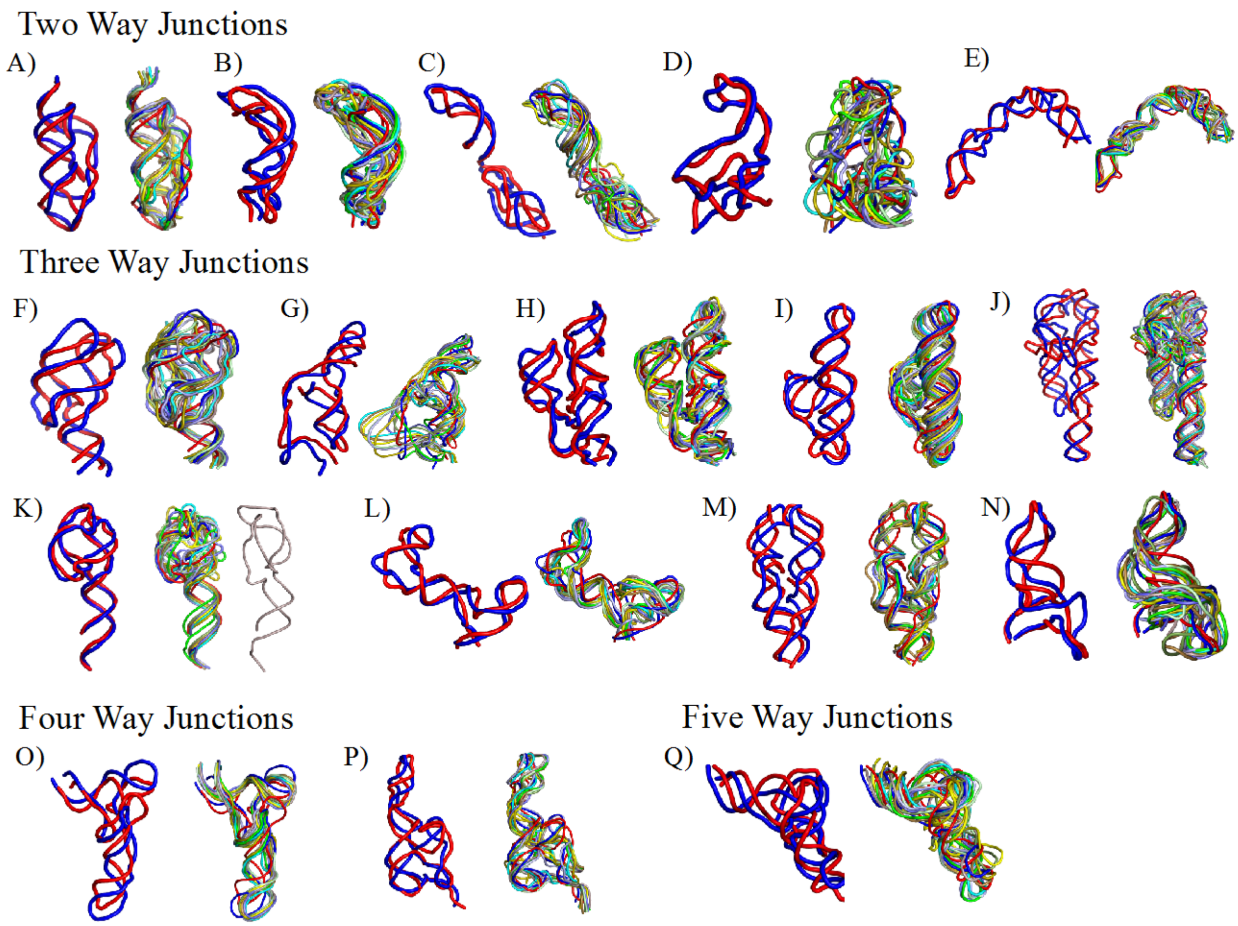

Figure 3.5 (A) through (Q) are the superimpositions of X-ray crystal or nuclear magnetic resonance (NMR) backbone structures with those of the RS3D structures. The X-ray/NMR crystal structures and the RS3D structures with the lowest root-mean-square deviation (rmsd) are colored in red and blue, respectively. On the right in each panel are the superimpositions of the top 10 structures with the lowest rmsd. Structures of 17 RNA species were superimposed.

3.2.6 Usage of the INCITE and ALCC Hours

The INCITE 2016 program allocated 3.5 billion core-hours on Mira. The allocation usage is shown in Figure 3.6. Of the 34 INCITE projects, 25 projects used more than 90 percent of their allocation. Twenty-one projects used their entire allocation (or more), including four projects using over 150 percent. These projects used the extra core-hours to achieve additional milestones. The overuse of Mira was made possible through the use of the backfill queue (low priority) and an "overburn" policy that permitted projects to continue running capability-sized jobs after their allocation was completely exhausted. The duration of the "overburn" policy was chosen to run from July 1 through November 30, 2016.

Of the remaining 13 projects with outstanding allocations, 10 projects used more than 50 percent of their time and of those 10, 5 projects used more than 75 percent of their time. Three projects used less than 50 percent of their allocation. The PI of one of these projects was on disability leave. The other two projects waited until too late in the year to complete their campaigns. A total of 3.9 billion core-hours were delivered to INCITE almost entirely on Mira. The contribution from Cetus to the total INCITE hours delivered was negligible (0.000092 percent).

For the 2015–2016 ALCC year, 26 projects had allocations on Mira for a total of 1.7 billion core-hours. The allocation usage is shown in Figure 3.7. Eleven of these projects used 90 percent or more of their allocations, including one project which also used Cetus for production runs that were not easily accommodated on Mira. Cetus usage accounted for about 9 percent of their usage and less than 0.3 percent of the total ALCC 2015–2016 usage. Of the remaining 15 projects, 7 used 75 percent of their allocation or more and only 4 projects used less than 50 percent of their allocation. Of those four, one project suffered a rescoping of project milestones owing to project understaffing. A second project encountered an issue with its methodology. Although ALCF staff helped the team resolve the problem, it was too late, and they ran out of time at the end. Another project started computing very late, and they were still able to consume close to 40 percent of the allocation. The last project was an off-cycle project (awarded March/April 2016). The team did not manage to get started during the 2015–2016 ALCC cycle, but their award continued into 2016–2017, and they have now used 30 percent of the allocation.

The 2016–2017 ALCC year is approximately halfway through its allocation cycle. So far, 27 projects have received allocations of 1.7 billion core-hours. The projects have used a total of 443.4 million core-hours from July 1, 2016, through December 31, 2016. The allocation usage is shown in Figure 3.8. One of the projects has already used its allocation and is running in the backfill queue, while another eight projects have used 50 percent or more of their allocations. To date, no projects are using Cetus for runs that are not easily accommodated on Mira.

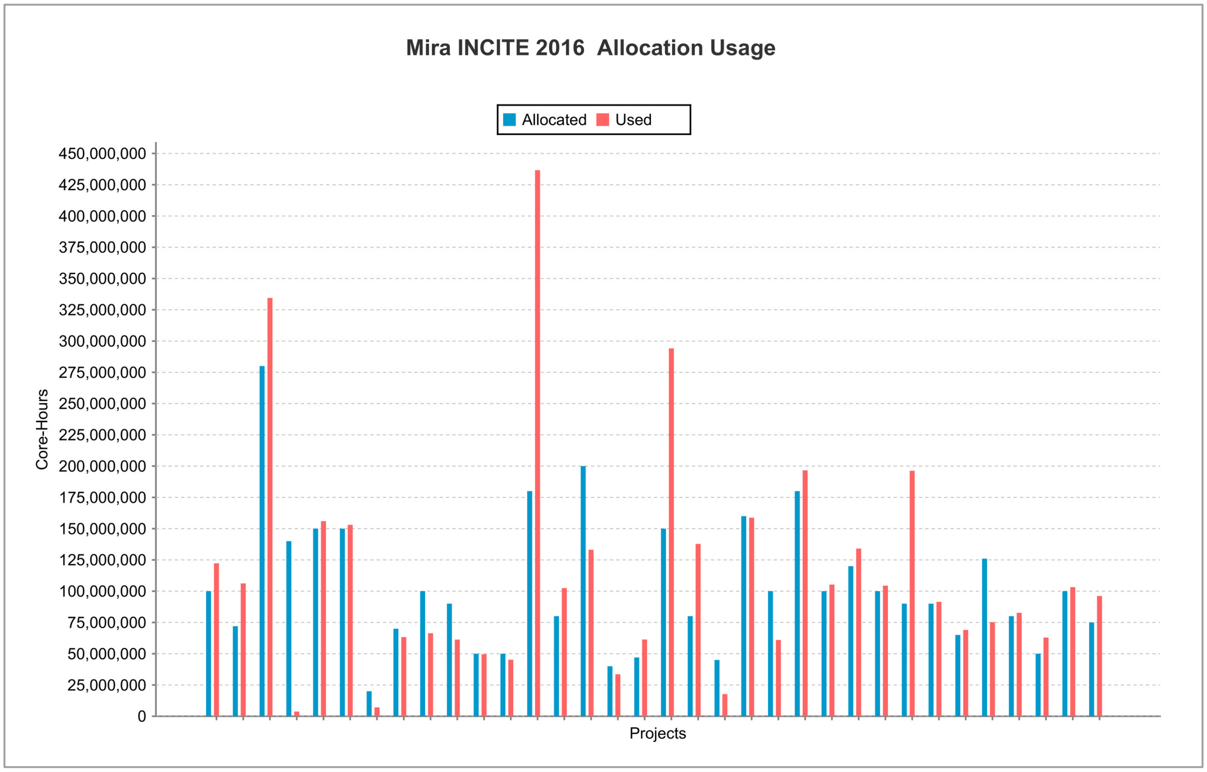

Figure 3.6 Mira INCITE 2016 Allocation Usage

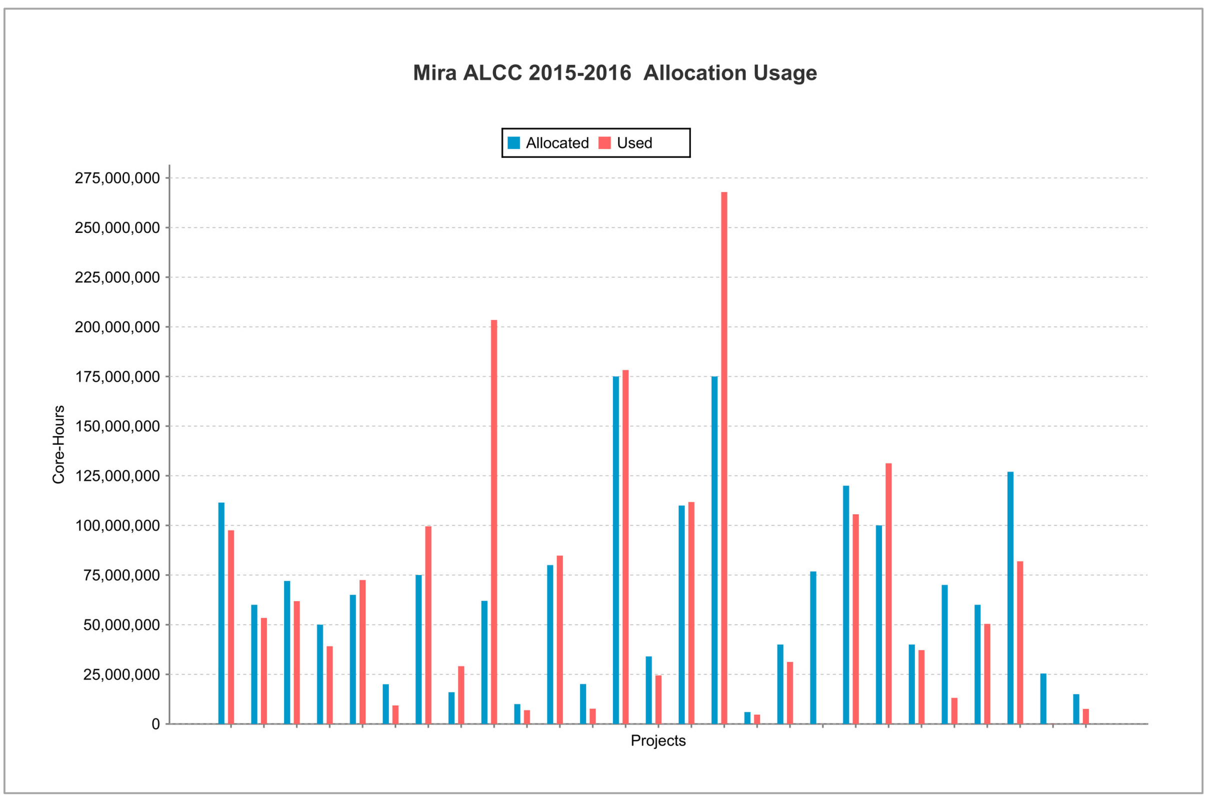

Figure 3.7 Mira ALCC 2015–2016 Allocation Usage

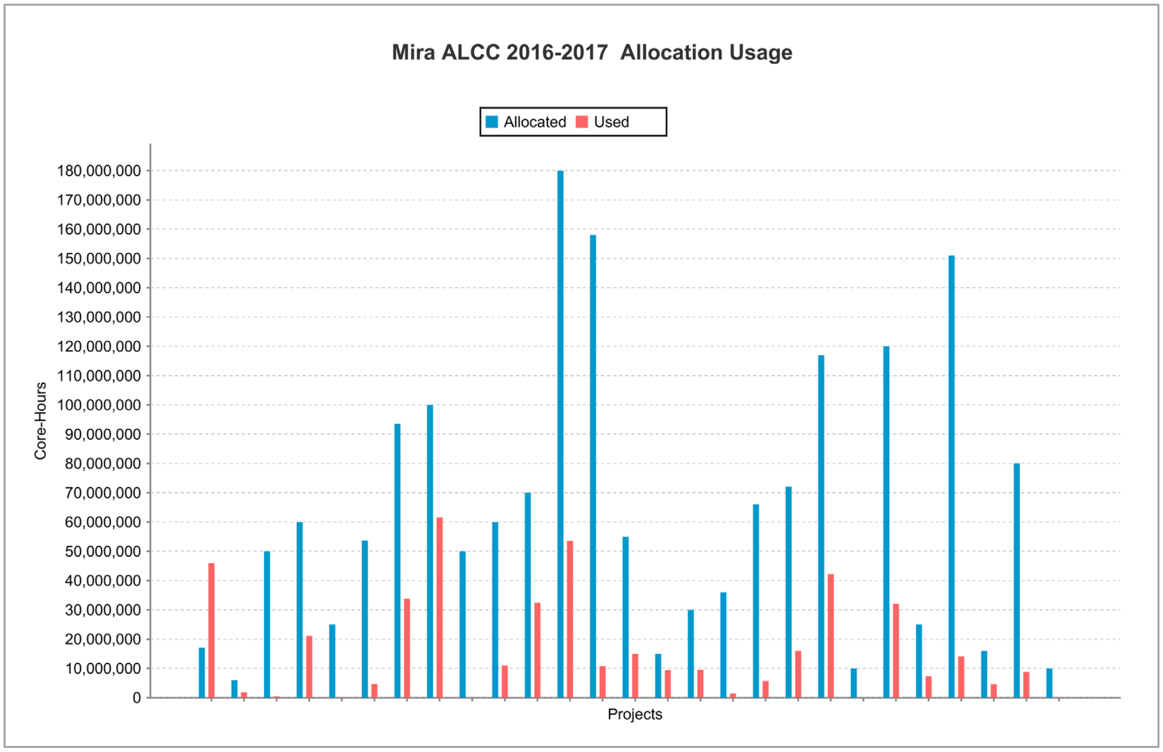

Figure 3.8 Mira ALCC 2016–2017 Allocation Usage

3.3 Allocation of Facility Director's Reserve Computer Time

In this section we are interested in the strategic rationale behind use of DD time. The Facility should describe how the Director's Discretionary reserve is allocated and list the awarded projects, showing the PI name, organization, hours awarded, and project title.

The Director's Reserve, or Director's Discretionary (DD) program, serves members of the HPC community who are interested in testing science and applications on leadership-class resources. Projects are allocated in four categories:

1) INCITE or ALCC proposal preparation
2) Code support and/or development
3) Strategic science
4) Internal/support

INCITE and ALCC proposal preparation allocations are offered for projects that are targeting submission of an ALCC or INCITE proposal. These projects can involve short-term preparation (e.g., a run of scaling tests for their computational readiness) or longer-term development and testing.

Code support and/or development allocations are used by teams porting and optimizing codes or projects developing new capabilities. This category includes the development, testing, and runs required for competitions such as the Gordon Bell Prize. Projects in this category have been responsible for bringing new capabilities to ALCF.

ALCF also allocates time to projects that might still be some time away from proposing for and receiving an INCITE award, or that offer a "strategic science" problem worth pursuing. Examples include supporting projects from DOE's Scientific Discovery through Advanced Computing (SciDAC) program, industry research efforts, and emerging use cases, such as coupling experimental and computing facilities.

Internal/support projects are devoted to supporting the ALCF mission. ALCF does not reserve core-hours for division activities. All activities come out of the DD allocation pool. This category regularly includes projects that help staff support the users and maintain the system, such as diagnostics and testing of tools and applications.

Allocations are requested through the ALCF website and are reviewed by the Allocations Committee (which includes representatives from Operations, User Experience, and the Catalyst teams). The committee collects additional input from ALCF staff, where appropriate. Allocations are reviewed on their readiness to use the resources and their goals for the allocations and are awarded time on a quarterly basis. The DD allocation pool is under great demand, and often the requested amount cannot be accommodated.

Table 3.2 shows the number of projects and total time allocated in the DD program on Mira during 2016. By its very nature, the DD program is amenable to over-allocation because often time is left unused; however, it should be noted that these totals do not represent open allocations for the entire calendar year. A project might have a 1-million core-hour allocation that only persists for three months, but that 1-million core-hour allocation is counted entirely in the annual total core-hour number. Projects are not guaranteed the allocated time; rather, the time is provided on a first-come, first-served basis. DD projects run at a lower priority than INCITE or ALCC projects, which reduces the amount of time available for their use. Exceptions are made for some internal projects that support acceptance of new hardware or support of users, which is in line with the ALCF core mission.

Table 3.2 DD Time Allocated and Used on Mira, 2016

Projects	Mira
Allocated Core-Hours	1.0B
Used Core-Hours	365.5M [a]

[a] Usage includes 21.2M core-hours from Cetus production jobs.

A list of the CY 2016 DD projects, including title, PI, institution, and hours allocated, is provided in Appendix A.

Figure 3.9 provides a breakdown of the CY 2016 DD allocations by domain.

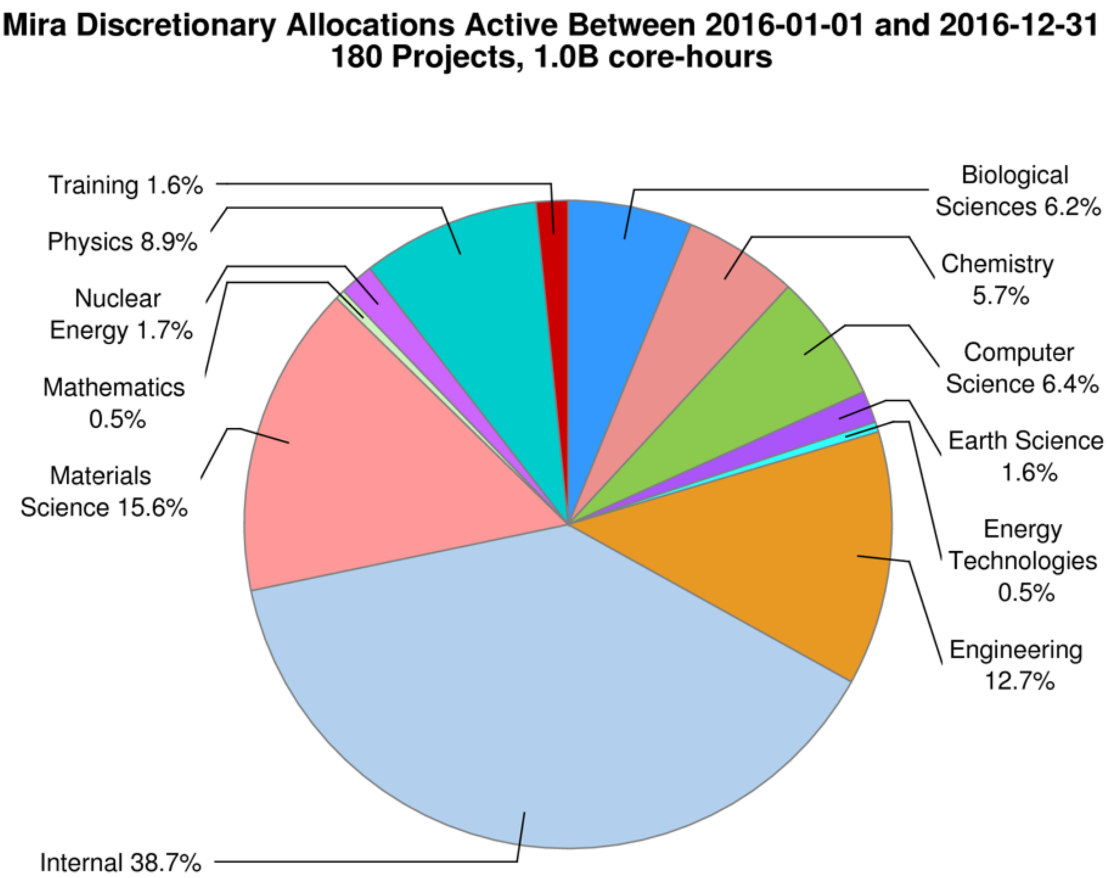

Figure 3.9 CY 2016 DD Allocations by Domain

3.4 Joint Laboratory for System Evaluation (JLSE)

The Joint Laboratory for System Evaluation (JLSE) is a joint collaboration between the MCS division and ALCF with the aim of evaluating future HPC platforms. The JLSE supported more than 25 projects and 75 users in 2016. Users were provided access to several new systems including a 10-node Xeon Phi (KNL) OmniPath cluster; a Cray Urika-GX analytics cluster with 32 Haswell nodes; an IBM S822LC Firestone system with a Power8 CPU and an Nvidia K80 Tesla GPU; a Broadwell Xeon cluster; and ARM processor from Applied Micro; and a Lustre testbed. The projects ranged from application portability to software development to tools and compiler development for an ALCF Early Science Project. Highlights from a few of these projects are included below:

ALCF Early Science: Twelve projects involving a total of 42 users were provided access to the Xeon Phi cluster at JLSE. The projects worked on single node optimization focusing on memory modes and vectorization. All 12 application codes were readied to run on Theta before the commencement of acceptance.

Argo: Argo is a new exascale operating system and runtime system designed to support extreme-scale scientific computation. The Argo project was given access to 90 dedicated nodes on an Intel cluster to prototype the GlobalOS distributed resource management across HPC systems. Performance evaluation of NodeOS was done using dedicated NVRAM hardware and Xeon Phi KNL boxes. Argobots/BOLT development and optimization of a lightweight low-level threading and task framework for OpenMP and other programming models (Cilk, Quark, Charm++) was also done at JLSE.

LLVM: LLVM compiler development is carried out on JLSE Power systems. The instruction set architecture (ISA) for these systems is the same as for BlueGene/Q with the only difference being in vectorization. LLVM and CLANG builds are carried out on the Xeon systems for QA purposes. These builds take 10 minutes (compared to hours on a laptop).

MPI: A number of MPICH improvements were tested on JLSE systems, including memory scalability of MPI communicators by exploiting regular patterns in rack-address mapping; enhanced threading support through locking optimizations; and communication-aware thread scheduling.

TAU: Argonne's Tuning and Analysis Utilities (TAU) performance system was ported to various ISA and architectures at JLSE. These include ARM64 (Applied Micro's X-Gene), Xeon Phi's (both KNL and KNC), and Power systems (Power7 and Power8). Support for OpenMP profiling via OpenMP tools interface (OMPT) and vectorization intensity were validated on JLSE systems.

Vegas: Application testing and development for Vegas was done on 44 CPU Broadwell nodes at JLSE. Vegas is written in C with MPI + OpenMP and uses domain decomposition on the geometry of the reactor. Vegas uses a new hybrid meta object compiler (MOC) and Monte Carlo method known as The Random Ray Method (TRRM). TRRM uses a stochastic quadrature to provide on-the-fly integration without any storage of quadrature information (thereby saving up to petabytes of memory). Use of JLSE resources was highly beneficial to test the multicore scalability of the application. Additional testing on Xeon Phi architectures was also useful to help assess the new algorithm's performance on manycore architectures.

The JLSE also presented monthly deep dive calls with Intel and Nvidia on topics ranging from chip architecture to algorithms to various software products. These calls were useful for understanding vendor technology roadmaps and products.

Conclusion

The ALCF continues to enable scientific achievements, consistent with DOE's strategic goals for scientific breakthroughs and foundations of science, through projects carried out on facility machines. Researchers participating in projects using ALCF resources published 254 papers in CY 2016. ALCF projects have had success in a variety of fields, using many different computational approaches. They have been able to reach their scientific goals and successfully use their allocations. A number of the projects and PIs have subsequently received awards or have been recognized as achieving significant accomplishments in their fields.

The ALCF delivered the following core-hours to the allocation programs in CY 2016: 3.9 billion to INCITE, 1.7 billion to ALCC, and 365.5 million to DD. The DD Reserve has been used not only to develop INCITE and ALCC proposals but also to conduct real science of strategic importance and to drive development and scaling of key INCITE and ALCC science applications. Excellent ALCF support and solid, high-performing ALCF resources have enabled INCITE and ALCC projects to run simulations efficiently on HPC machines and to achieve science goals that could not otherwise have been reached.

This page intentionally left blank.

Section 4. Innovation

Have innovations been implemented that have improved the facility's operations? This includes innovations adopted from, recommended to, or adopted by other Facilities.

ALCF Response

Listed below are the innovations and best practices carried out at ALCF during CY 2016. ALCF innovations and best practices have helped to prepare for future systems, have enabled more efficient operations, and have strengthened collaboration and engagement, both across ASCR and beyond.

4.1 Innovations

4.1.1 ALCF Data Science Program (ADSP)

The ALCF Data Science Program (ADSP) is part of ALCF's DD allocation program and is targeted at "big data" science problems that require the scale and performance of leadership computing resources. The goal of the program is to help explore and improve a variety of computational methods that will help enable data-driven discoveries across all scientific disciplines. The new initiative is the first program in the DOE computing facilities to tackle this new area. ADSP projects are expected to cover a wide variety of application domains that span experimental and observational sciences. The projects will focus on data science techniques including but not limited to uncertainty quantification, statistics, machine learning, deep learning, databases, pattern recognition, image processing, graph analytics, data mining, real-time data analysis, and complex and interactive workflows.

The ADSP awards are for two years and will be renewed annually. The program will have a regular call for proposals. The first ADSP call for proposals was issued in the spring of 2016 with a deadline of June 3, 2016. The ADSP projects will be categorized as either "data science projects," which will have a specific science goal, or "software technology projects," which will be focused on implementation of the specific technology required to support data science, including complex workflows. The ADSP projects selected in 2016 span experimental and computational sciences and range from modeling the brain to discovering new materials for solar-powered windows to simulating collision events at the Large Hadron Collider (LHC). Through ADSP, these projects will help the ALCF explore and improve a variety of computational methods that will enable data-driven discoveries across all scientific disciplines. The projects will be split into two tiers and will have access to the various ALCF and Joint Laboratory for System Evaluation (JLSE) computing resources as well as staff. Tier 1 projects will also be supported in part by postdoctoral researchers. ADSP teams also will be among the first to access Theta. Additional resources for this work include ALCF's Mira and Cooley systems and Sage, a Cray Urika-GX analytics cluster, at Argonne's JLSE. The next call is expected in spring 2017.

4.1.2 Cobalt Data Staging

Work has been done to add data-staging capabilities to Cobalt. This new functionality allows Cobalt to stage user-specified files to ALCF storage from external facilities using Globus Online to broker the transfer in a secure manner. This ability has been targeted for use in a collaboration with Argonne's Advanced Photon Source facility to enable its researchers to streamline their workflows and use ALCF resources more efficiently in their research. In addition, this feature is a critical first step toward an additional group of data-centric features in Cobalt that are targeted toward the Aurora system, notably handling nonvolatile node-local storage and burst-buffer capabilities.

4.1.3 RAN (RAM Access Network)

The RAN (RAM Access Network) project is intended to enable research into deep memory hierarchies and potentially different cluster memory configurations. ALCF has a 5.5 TB pool of RAM available over the InfiniBand (IB) network, with 14 FDR IB (Fourteen Data Rate InfiniBand) cards available for bandwidth on the pool side and single FDR cards on each node of the Cooley visualization cluster. This configuration enables bandwidths that are on par with those expected future systems. The latencies are typically within 1 microsecond of the raw IB latencies and extremely consistent in our testing so far. This low jitter should enable consistent performance, and for stream, we have achieved 50% of the local memory bandwidth capacity with only 128 KB of local buffer for read ahead. The memory pool is disk backed, so it can be considered persistent memory, which makes it a reasonable proxy for future NVRAM applications. We have a number of interfaces available to the memory pool including malloc, mmap, memkind, block, and QEMU/KVM/RHEV, as well as a direct C API that enables maximum performance by bypassing the kernel. Collaborators have also implemented a pmem-io persistent memory interface, as well as a simple MPI-IO interface, which was used in a legacy application that could not be modified, and enabled a 600–800× speedup when compared to writing to their Lustre file system. Application work in progress includes modifying the Argonne-written vl3 visualization tool to use this memory pool to enable streaming through the GPU (a paper on this work was presented at the ISAV [In Situ Infrastructures for Enabling Extreme-scale Analysis and Visualization] workshop at SC16), and a DDI interface for the GAMESS application.

4.1.4 Operational Improvements

4.1.4.1 Leadership Class Operational Data Analysis System

The ALCF's operational data analysis system was designed to track usage, availability, hardware state changes, reservations, jobs, or any event that happens over time that uses a compute resource. The system was created because the previous system, built for the Blue Gene-class supercomputers, was highly integrated and required manual validation of events. The new system is very generic and works up to very large size systems, which the previous system could not easily accommodate. It also enables the automatic validation of data and verification of results. This system will be the core of the data analysis systems for ALCF's DOE metrics and hardware/job failure analysis. It is currently being used to generate availability, usage, and utilization data for Theta in a pre-production state. The system was built for processing ALCF metrics much more quickly and for providing a robust API to the data. Processing a month's worth of data for Theta takes around two minutes. ALCF staff members have only begun to explore and leverage the power of this system.

The system is broken into three distinct pieces: the extract, transform, load (ETL); the analysis core; and business logic. The job of ETL is to load data into the system. The ETL is also the place where any differences between the IBM Blue Gene and the Intel Cray are resolved. This is called System Logic. The analysis core itself contains the database and a series of functions that allow manipulation of the data so that it can be processed by the business logic. The business logic is site specific and is kept separate from the rest of the system so that it only handles the specifics for that site. The business logic might specify how and what layers of events to include and how they are processed and joined together. The efficiencies afforded by the bit mask and the joining of all these data allow for analyses that would have been too time consuming or complicated in the past.

The system was originally built for an IBM Blue Gene/Q but was later ported to an Intel Cray machine. During that port, the appropriate abstractions were considered to make future ports to new machines or architectures easier. The ALCF plans to open source this system in the future. There are also plans to automate more of the manual systems to further improve facility efficiency and accuracy.

4.1.4.2 Bloodhound

Bloodhound is designed to monitor multiple high-performance computing appliances and servers and report on any hardware or software changes. These changes can include hardware faults, hardware properties leaving tolerable levels, or software changes such as configuration properties changing. All of these state changes are stored in a back-end database for future reference and analysis. In addition to state changes, Bloodhound specifically tracks hardware replacements, providing a detailed list of the hardware that has touched the system and when. Bloodhound can monitor any system so long as it knows how to “talk” with a given system; currently supported systems include DDN SFA12K-20E and SFA12K-40 and IBM ESS (elastic storage server).

This tool was created because the storage team needed a way to monitor all of the storage appliances in a consistent and organized manner such that the storage appliances would be easily expandable to new future hardware. Previous methods included email alerts sent by each box, that resulted in noisy and often false positives, or scripts that focused only on one system type or a small subset of variables. A major need was to acquire the ability to track the physical hardware and to know when something was installed, how long it has been in the system, and when it was removed. ALCF also wanted the ability to look back at past system states in case there was a fault and it became necessary to know what the system looked like at that time.

Bloodhound is currently monitoring the DDN SFA (DataDirect Network storage fusion architecture) storage and IBM ESS storage for hardware and software changes. On a daily basis, the storage team receives a report on any change that may need to be addressed. The major types of changes seen are hardware faults, and these can be addressed quickly, as Bloodhound provides the exact location of the faulted component. Another benefit to using Bloodhound is the capability it provides to store the system state data off of the monitored system. Should that system have a catastrophic failure, a portion of its state will be available to help expedite the recovery. ALCF has already seen the benefits of Bloodhound beyond the hardware monitoring; it caught a sweeping change to a few DDN parameters that could have greatly affected performance. The cause of the change was found, stopped, and the parameters quickly reset to their previous values.

4.1.5 BG/Q Provider Work for OFI

The ALCF and Intel have developed Libfabric Open Fabric Interface (OFI) for the Blue Gene/Q. First, this interface will provide a messaging-layer alternative to PAMI on Blue Gene /Q that enables a fully functional implementation of MPI-3 for Mira and allows ALCF staff to investigate MPI performance improvements. Second, use of this OFI will aid in Libfabric OFI and MPI development for Aurora by allowing developers to leverage a currently accessible and highly scalable architecture (Mira) to explore and develop certain aspects of Libfabric OFI for performance and scalability. This interface is open-sourced and contributed to the OFI Working Group GitHub site.

4.1.6 Project Allocation Analysis

The ALCF sends weekly project status reports to all INCITE and ALCC Principal Investigators. These reports now include a project-specific usage table that presents the number of days and relative job sizes required by the project to use 100% of its allocation before it expires. Figure 4.1 shows an example of this table for the GMSeismicSim INCITE project. Based on this table, it was clear that continuing to run 512- and 1024-size jobs, as they primarily had been, would not allow them to reach their goal. Working with ALCF staff member Tom Uram, the team members changed the way they submit jobs, leading to capability jobs accounting for the majority of their usage, as seen in Figure 4.2

Job Size	Full days needed at Job Size	Days left in allocation	Number of past jobs run at Job Size
512	358	76	934
1024	179	76	34
2048	89	76	16
4096	45	76	12
8192	22	76	14
16384	11	76	10
32768	6	76	8

Figure 4.1 Example usage table showing full days needed at job size and days remaining in current allocation for GMSeismicSim project.

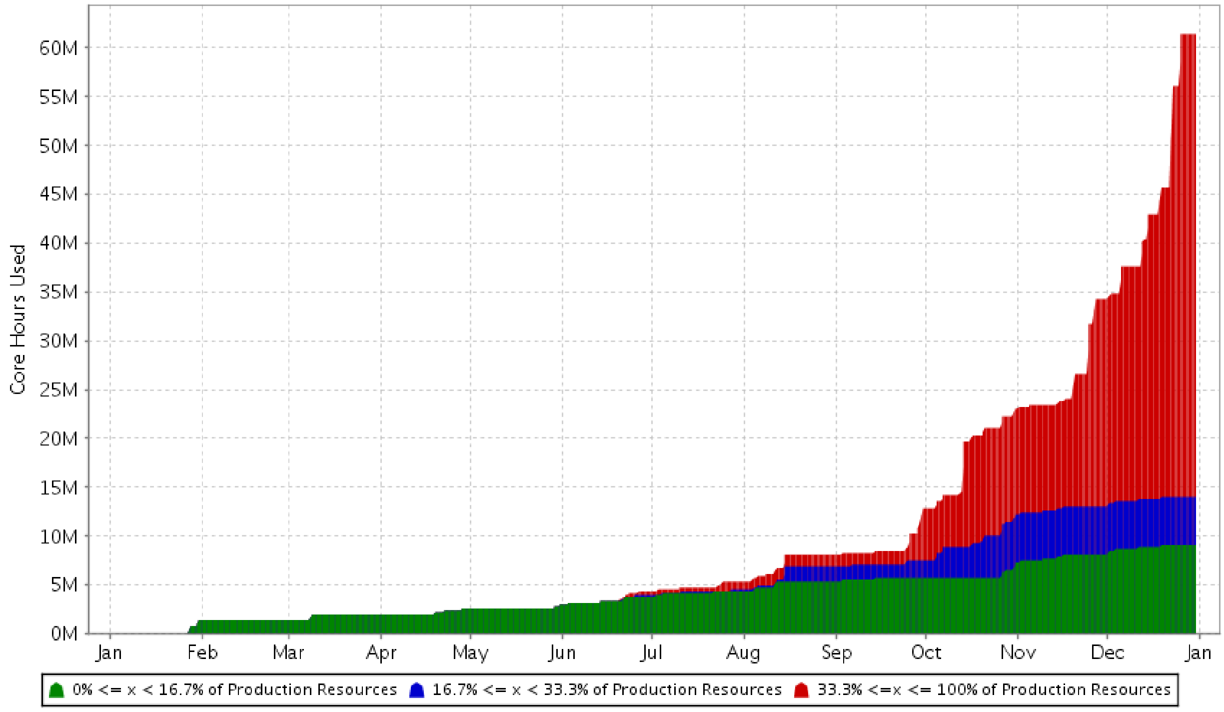

Figure 4.2 Mira usage graph for 2016 INCITE project GMSeismicSim.

4.2 Best Practices

4.2.1 Best Practice: Performance and Library Tracking

The more we understand about how ALCF resources are being used, the more informed decisions we can make about where to focus effort to best support our users. To that end, ALCF has developed and employed a number of tools for automatically monitoring various aspects of Mira, such as Darshan for investigating I/O behavior. In 2015, the experimental deployment of TrackLib took place for collecting data on compiler and library usage, and AutoPerf for the automatic collection of application performance data on Mira. In 2016, both of these tools went into production on Mira, and data collected by them are now incorporated into the ALCF data warehouse.

Since the deployment of AutoPerf on Mira, all codes compiled and run on the system have automatically incorporated the AutoPerf library, except for situations where the use of AutoPerf would conflict with other tools, in which case the conflict is detected and AutoPerf is automatically removed. Over the subsequent period of AutoPerf operation, no AutoPerf-related issues or performance impact on applications has been found or reported and no job failures have been attributed to the library. During this time, AutoPerf has collected data for over 150,000 jobs that have ranged in size from 1 to 1 million MPI ranks and in duration from micro-seconds to 24 hours. To date, the collected AutoPerf data have been utilized, for example, to examine and characterize typical MPI usage patterns for applications run on Mira to guide and inform the Aurora non-recurring engineering (NRE) effort on optimized MPI collectives.

TrackLib enables tracking of libraries by recording the input files used to build executables. Later, at runtime, executables are scanned, and runtime information is logged. The data can then be analyzed to help guide decisions on what software stack components to support. TrackLib data collected on Mira have been used to inform the Exascale Computing Project in the application and software decision process.

4.2.2 Best Practice: Mira's Service Node Data Protection Strategy

Mira's control system has been set up to provide high availability (HA). The control system is hosted on two Linux systems in a Linux-HA configuration, allowing system software components to fail over should one of the nodes become unusable. Furthermore, the database instances on which the control system relies are themselves in a high-availability/disaster recovery pair. Loss of one node automatically initiates standby database activation, after which control system components can be failed over. This configuration allows ALCF to recover from a service node failure after only a short downtime.

Furthermore, data protection and disaster recovery are provided for these components. Full backups are taken on a bi-monthly basis, along with weekly incremental, daily delta, and transaction log backups. In the event that the database becomes corrupted, a backup taken prior to the event can be restored, and the transaction logs can be used to roll forward to any point in time during that day. Consequently, recovery from any type of data integrity event can be accomplished within a few hours and would result in the loss of only data written *following* the event. Without this capability, it would be necessary to reinitialize the database and perform an import of the old data, a process that could easily take days to complete.

4.2.3 Best Practice: Management of Legal Documents

ALCF now uses a single point of contact for the collection and management of all legal documents related to NDAs, open-source agreements, and license agreements. These documents are stored in a repository that is available to all ALCF staff to verify the facility's agreement with specific institutions and specific restrictions on information sharing.

4.2.4 Best Practice: Use of Standard Software for Proposal Management

For the Aurora Early Science Program, ALCF used, for the first time, the publicly available EasyChair Internet-based software to manage submissions, interact with submitters and reviewers, and manage the workflow. This reliance on a commonly used submission management system helped make the submission, review, and decision process operate more smoothly.

4.2.5 Best Practice: Standards Committees

A necessary part to the ALCF mission is to lead the HPC community into the future. Serving on standards committees is a best practice in this mission.

For C++17, ALCF has helped direct the design of, and pushed for, parallel standard-library algorithms, lambda functions that capture the current object by value (important for the construction of parallelism abstraction libraries), and other features. The "inline variables" feature, which will make it easier to construct header-file-only libraries, was a feature on which ALCF was a primary co-author. For future revisions of the standard, the ALCF staff are also working on pointer-aliasing attributes to enable higher-performance code.

ALCF participates in the weekly OpenMP ARB (Architectural Review Board) and Language Committee teleconferences, and ALCF staff have attended the Spring face-to-face meeting. (Members of the MCS's Message Passing Interface Chameleon (MPICH) team represented Argonne at the Fall face-to-face meeting.) Following the release of OpenMP 4.5 in November of 2015, the primary committee activity for 2016 has been to publish previews of functionality for the OpenMP 5.0 release: Technical Report 4 (including OMPT, the standardized interface for tools such as profilers) and Technical Report 5 (Memory Management Support, to support both multiple classes of memory and different spaces such as those of accelerators).

4.2.6 Best Practice: Collaboration with ASCR and Other Facilities

The ALCF is part of a family of supercomputing facilities within the DOE. These facilities should not operate in isolation, and a best practice is to collaborate and share expertise across facility staff. Examples of this collaboration include the following:

- Share experience with new KNL systems. ALCF and NERSC have had conference calls on application development to share their experiences on the Early Science Program (ESP) and NERSC Exascale Science Applications Program (NESAP) projects. They also had a day-long, face-to-face meeting on June 13 of this year at Argonne to discuss a variety of topics from application code optimization, system administration, data management, burst buffer, etc. ALCF also takes an active part in biweekly conference calls organized by the Tri Labs to share experiences on KNL.
- Staff exchanges for hands on learning. An operations staff member [Ben Allen] visited LANL during its Cray XC40 SMW upgrade to learn first-hand how they perform upgrades and to observe their workflow. He learned many "gotchas" to look out for and many best practices. The various detailed steps covered updating OS images, how to configure and apply patchsets, the boot process, and various monitoring and metrics. A few weeks later, LANL staff member Paul Peltz spent time with the ALCF operations team discussing various topics of interest to both facilities. These meetings covered topics such as security, logging, and configuration management with Ansible as the primary topic; provisioning; GPFS; Lustre tuning; user environment; and networking. ALCF and LANL found both visits to be immensely useful.
- Collaboration on training. As mentioned, collaboration between the ASCR Computing Facilities and the ASCR Research IDEAS project resulted in a very well-attended webinar series covering introductory supercomputing topics.

This page intentionally left blank.

Section 5. Risk Management

Is the Facility effectively managing risk?

ALCF Response

ALCF has clearly demonstrated successful risk management in the past year for both project and operation risks. The risk management strategy is documented in the ALCF Risk Management Plan (RMP), which is reviewed and updated regularly to incorporate new ideas and best practices from other facilities. Risk management is a part of ALCF culture, and the RMP processes have been incorporated into both normal operations and all projects, such as the ALCF-3 project launched in CY 2013. Risks (proposed, open, and retired) are tracked, along with their triggers and mitigations (proposed, in progress, and completed), in a risk register managed by two Risk Co-Managers. All risk ratings in this report are post-mitigation ratings. ALCF currently has 38 open risks, with three high operational risks: (1) Funding/Budget Shortfalls, which is managed by careful planning with the DOE program office and the continuation of austerity measures as necessary; (2) Staff Recruitment Challenges, which is managed by ongoing recruiting and retasking current staff as needed; and (3) Catastrophic Failure of Home File System, which is managed by several types of file backups. The major risks tracked for the past year are listed in Section 5.1, with the risks that occurred and the mitigations for those risks described in greater detail in Table 5.1 and Section 5.2, respectively, along with new or recharacterized risks (Section 5.5, Table 5.3) and retired risks (Section 5.4, Table 5.2), as well as the major risks that will be tracked in CY 2017 (Section 5.6, Table 5.4).

Discuss how the Facility uses its RMP in day-to-day operations, how often the RMP is reviewed or consulted, and what happens when a risk occurs. For this review the focus is on Operational risks, not Project risks.

The Facility should highlight various risks to include:

- *Major risks that were tracked for the review year;*
- *Any risks that occurred and the effectiveness of their mitigations;*
- *A discussion of risks that were retired during the current year;*
- *The mechanism used to track risks and trigger warnings;*
- *Any new or recharacterized risks since the last review; and*
- *The major risks that will be tracked in the next year, with mitigations as appropriate.*

Note: *This is a high level look at the risks, not a deep dive into the risk registry.*

5.1 ALCF Risk Management

ALCF uses the documented risk management processes, first implemented in June 2006 and outlined in its RMP, for both operations and project risk management. ALCF reviews and updates the plan annually. The plan is also updated as needed during the year to reflect changes and to incorporate new risk management techniques as they are adopted by the facility. The RMP is consulted at all monthly and individual risk meetings. Details of the RMP, including the attributes of each risk managed by ALCF, have been described in past reports and will not be discussed further here. Risks are tracked using a secure shared cloud-based storage system, and risk forms and the risk register are formatted using Excel. Risk owners continuously monitor the risks they own and submit monthly reports on all risks.

5.1.1 Continuation of the ALCF-3 and ALCF-Lithium Projects

The ALCF-3 project — procuring and deploying the next ALCF supercomputer — continued in CY 2016. Risk Register managers continue to maintain a project risk register and track a set of detailed risks. Risk mitigation costs on the project side are developed using a bottom-up cost analysis, then are input to the commercial project risk analysis tool Oracle Primavera Risk Analysis (OPRA) to set the contingency pool utilizing the OPRA integration with the Primavera project management tool. These risks are not included in the risk numbers covered in this document and are not discussed further.

The ALCF-Lithium project — a special project to procure and deploy the Theta supercomputer — continued in CY 2016. The project also maintained its own risk register, which was tracked and managed utilizing the same tools as the ALCF-3 project. With the acceptance of Theta at the end of September 2016, that risk register was closed out. All risks that remained in the Theta risk register were evaluated and verified to be covered under the current set of ALCF steady-state risks.

5.1.2 Collective Risk Review and Updating

As an addition to the individual monitoring of each steady-state risk by the risk owner, a collective review of all risks was performed in CY 2016. ALCF management, the ALCF Risk Review Board, and the Risk Register managers worked with the risk owners to evaluate the risks as a whole. The results included identification of several risks that could be consolidated, improved clarity and consistency in risk wording, and some scoring adjustments after considering the broader picture of ALCF risk mitigation efforts and historic data on risk encounters.

5.1.3 Online Risk Reporting Form

An online steady-state risk reporting form was developed and launched in CY 2016 to streamline the process of monthly reporting by the risk owners and of maintaining written records of the monthly reports and monthly summaries of the reports. Risk owners receive an e-mail reminder at the beginning of the month to submit a report for the month just concluded on the risks that they own; when they click the link to the form, a form that is customized to display only their risks appears. An example form is shown in Figure 5.1 for the risks owned by the HPC Infrastructure team lead.

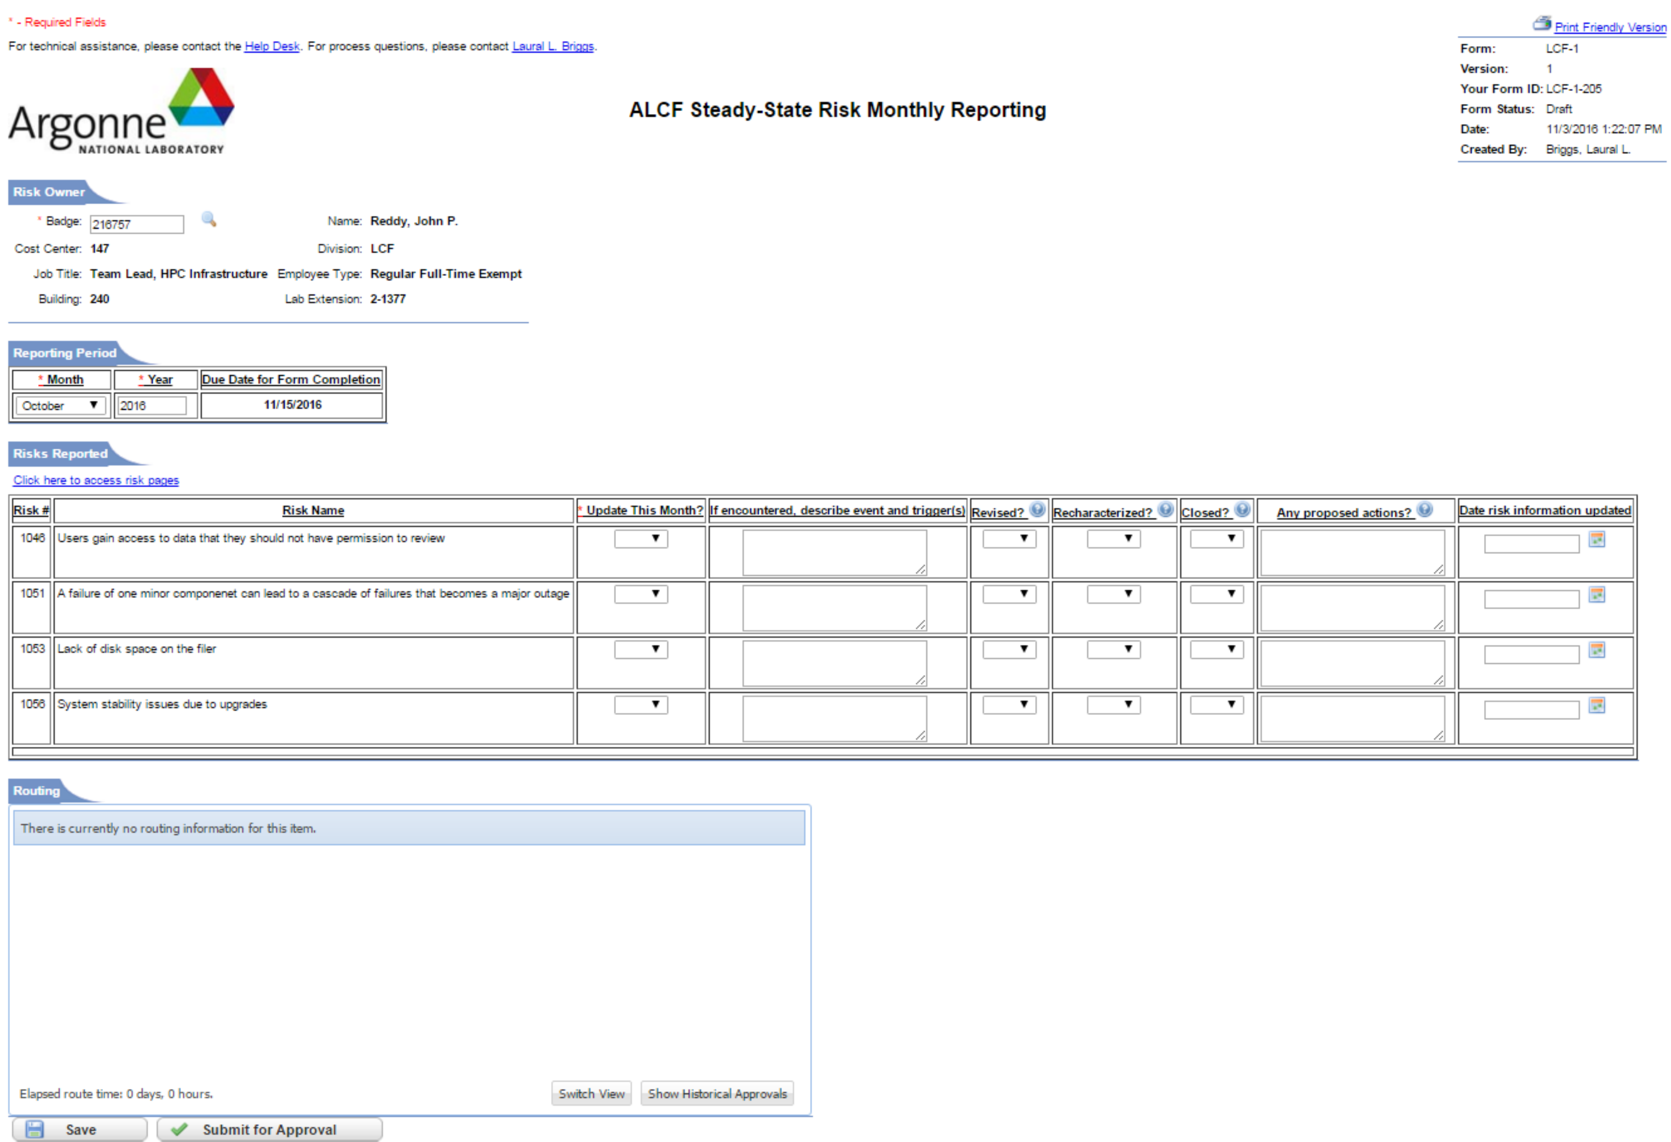

* - Required Fields

For technical assistance, please contact the Help Desk. For process questions, please contact Laural L. Briggs.

Print Friendly Version

Form:	LCF-1
Version:	1
Your Form ID:	LCF-1-205
Form Status:	Draft
Date:	11/3/2016 1:22:07 PM
Created By:	Briggs, Laural L.

Argonne NATIONAL LABORATORY

ALCF Steady-State Risk Monthly Reporting

Risk Owner

* Badge:	216757	Name:	Reddy, John P.
Cost Center:	147	Division:	LCF
Job Title:	Team Lead, HPC Infrastructure	Employee Type:	Regular Full-Time Exempt
Building:	240	Lab Extension:	2-1377

Reporting Period

* Month	* Year	Due Date for Form Completion
October	2016	11/15/2016

Risks Reported

Click here to access risk pages

Risk #	Risk Name	* Update This Month?	If encountered, describe event and trigger(s)	Revised?	Recharacterized?	Closed?	Any proposed actions?	Date risk information updated
1046	Users gain access to data that they should not have permission to review							
1051	A failure of one minor componenet can lead to a cascade of failures that becomes a major outage							
1053	Lack of disk space on the filer							
1056	System stability issues due to upgrades							

Routing

There is currently no routing information for this item.

Elapsed route time: 0 days, 0 hours.

Switch View | Show Historical Approvals

Save | Submit for Approval

Figure 5.1 New ALCF Online Steady-State Monthly Risk Reporting Form Example

If a risk was encountered during the reporting month, the owner can enter into the form all information describing the encounter and actions taken. If actions related to other risk updates (revision, recharacterization, closure) were taken, these can also be recorded through the form. There is also a link to the folder where copies of all risks are archived, so that the owner can access the risk statement directly from the form if desired when documenting a risk encounter. Once the form is completed, the owner can submit the report for approval. A report can also be updated online if new information becomes available concerning a risk that was encountered.

The online form submits the risk report to the chair of the ALCF Risk Review Board for approval. The Board chair can then use another online tool to generate a monthly summary of the risk reports in an Excel spreadsheet. Reports are archived in a secure shared cloud-based storage system.

5.1.4 ALCF Risk Review Board

ALCF employs a five-person Risk Review Board to serve in an advisory capacity to ALCF management. The board meets on a quarterly basis and makes recommendations to ALCF management regarding steady-state risk management issues. At each meeting, the board:

- Reviews proposed new risks and makes recommendations on adding a proposed risk to the steady-state ALCF risk register.
- Monitors open risks and, for each open risk, reviews any new information on the risk provided by the risk owner and/or the ALCF steady-state risk managers and:
  - Determines whether the risk needs to be recharacterized.
  - Considers whether the risk has been managed and should be closed.
  - Reviews the mitigation actions for the risk and considers whether any of the actions need updating.

5.1.5 Risk Management in Day-to-Day Operations

ALCF currently has 38 open risks in the facility operations risk register and uses the post-mitigated risk scoring to rank the risks. These risks include general facility risks (such as funding uncertainties, staffing issues, and safety concerns) and specific risks (such as system component failures, availability of resources, and cost of electricity). On the operations side, subject matter experts estimate risk mitigation costs and use them to inform management reserves.

In addition to formal monthly and individual risk meetings and the Risk Review Board quarterly meetings, ALCF has many informal risk discussions. Risks are identified and evaluated, and mitigation actions developed, for all changes that occur at the facility, from installing a new piece of hardware, to changing the scheduling policy, to upgrading software. If the risks identified are short-term or minor, they are not added to the registry. New significant risks identified during the activity planning are added to the registry and reviewed at the next risk meeting.

Other tools beyond the risk register are used for managing risks in day-to-day operations. An example is the use of Work Planning and Controls (WPCs) and Job Hazard Questionnaires (JHQs) to manage risks for activities where safety is a potential concern. WPCs are primarily used for any nonroutine work and are developed in consultation with safety and subject matter experts. JHQs are used for all staff and all contractors and cover all work, both routine and nonroutine. During planning meetings for nonroutine activities, staff members review the planned actions and evaluate possible safety concerns. If a potential risk is identified, detailed discussions with the safety experts are scheduled, and procedures for mitigating the risks are developed and then documented in the WPC. The WPC is then used during the activity to direct the work.

Beyond the operations of the machine, risk management is used in such diverse ways as in evaluating and managing INCITE and ALCC proposal risks (the risk of too few proposals, the risk of a lack of diversity across science domains, the risk of too few capability proposals, etc.), safety risks in staff offices, leasing risks, support risks (including the opportunity risk that electricity costs could be lower than budgeted), etc.

5.2 Major Risks Tracked for the Review Year

ALCF has experienced another eventful year as a result of transitioning Theta to operation during the last half of CY 2016 and the planned growth of both ALCF staff and budget in order to bring the facility to full strength. As such, ALCF continues to monitor a large number of major risks for the facility. No major risks were retired during CY 2016.

Six major operations risks were tracked for CY 2016, three with a risk rating of High and three with a risk rating of Moderate. Of these, three were encountered and managed. The six major operational risks are described in Table 5.1. All risk ratings shown are post-mitigation ratings. The risks are color-coded as follows:

- Red risks were encountered and remain Moderate or High risks.
- Orange risks were not encountered but remain Moderate or High risks.

Table 5.1 Major Risks Tracked for CY 2016

ID	Title	Encountered	Rating	Notes
25	Staff recruitment challenges	Yes	High	ALCF added 6 new hires overall this year, plus 1.5 internal transfers. ALCF continues to have staff available who can be re-tasked as needed. With ongoing budget uncertainties and difficulty competing with industry for new hires, staff hiring remains a concern.
1059	Funding/budget shortfalls	Yes	High	ALCF worked with the program office to plan a budget for handling the impact of a Continuing Resolution, new hires, and changes in the laboratory indirect expense rate. This risk remains a major concern as the facility moves forward with Theta and ALCF-3.
1049	Staff retention	Yes	Moderate	ALCF lost three staff members during CY 2016. Budget concerns at Argonne and the growth in high-paying industry jobs for system administrators and programmers with HPC expertise make staff retention in out years a continuing concern.
1054	Catastrophic failure of home file system	No	High	This is a risk that has a low probability but a very high technical impact if it were to occur. It is a reportable event that would fully shut down the machine. However, the file system could be restored within a day or two from the mirror of the home file system that is maintained.
1018	INCITE and ALCC users are not provided adequate support by ALCF	No	Moderate	ALCF staff is proactive about limiting the chance of encountering this risk by (1) frequently soliciting feedback from the user community about the service ALCF provides, and (2) actively managing the support expectations of the user community.
1044	Unable to meet INCITE Office of Management and Budget (OMB) metric	No	Moderate	Appropriate management by ALCF staff keeps the probability of this risk very low and the overall score moderate, despite a very high technical impact score that reflects a failed OMB metric if the risk is encountered.

5.3 Risks Encountered in the Review Year and Their Mitigations

The seven risks encountered during CY 2016 are discussed below, along with the risk owner, the risk probability and impacts, a description of the actual problem that occurred, and the management of the risk. The ratings of the risks encountered were as follows: 2 High, 1 Moderate, 2 Low, and 2 Very Low.

5.3.1 Funding/Budget Shortfalls

1059: Funding/Budget Shortfalls	
Risk Owner	Michael Papka
Probability	High
Impact	Cost: Very Low; Technical Scope: High
Risk Rating	High
Primary Management Strategies	Develop austerity measures. Work closely with DOE sponsors to manage expectations and scope. Plan carefully, in conjunction with program office, for handling Continuing Resolution, leasing costs, and hires. Forward-pay lease to reduce overall leasing costs.
Triggers	ASCR provides funding for budget that is less than planned. Information from DOE indicates a likely extended Continuing Resolution. Argonne laboratory management calls for austerity measures.

Description

The Office of Science might not fund the ALCF budget as planned, or could reduce the ALCF budget below previous funding levels. An extended or full-year Continuing Resolution could prevent ALCF from receiving planned funding. These scenarios could result in the inability to pay leases, contracts, and staff and to deploy future machines.

Evaluation

During the past year, the Funding/Budget Shortfalls risk was one of ALCF's three highest risks, and it was also one of the risks encountered. The facility was required to operate with moderate austerity measures during the early part of the year. ALCF plans for carryforward funds each year, with the intention of starting each fiscal year using carryforward funding from the previous fiscal year while waiting for the first allotment of current year funding to arrive. This approach was followed in FY 2016; however, the funding uncertainty was large enough to also require some budget reprioritizing of purchases and new hires.

Management

In conjunction with the DOE-ASCR Budget Deep Dive, ALCF prepared for a full-year Continuing Resolution and reduced budget scenarios. To assure that adequate funds were available to operate Mira and prepare for Theta and ALCF-3, ALCF continued moderate austerity measures to provide maximum flexibility for the coming fiscal year.

ALCF continues to closely monitor budget information for FY 2017 and beyond in case of a reduction in funds from the plan of record. Moderate austerity measures remain in place, with spending being prioritized, and these measures may be augmented, depending on the budget.

5.3.2 Staffing Recruitment and Retention Challenges

25: Staff Recruitment Challenges	
Risk Owner	Michael Papka
Probability	Moderate
Impact	Cost: Low; Technical Scope: High
Risk Rating	High
Primary Management Strategies	Evaluate possible additional recruiting avenues. Prioritize staffing needs. Adjust work planning. Retrain staff to meet ALCF needs. Retask staff as needed. Leave job postings active and open.
Triggers	Lack of response to job postings. Rejection of job offers. Staff turnover.

1049: Staff Retention	
Risk Owner	Michael Papka
Probability	Moderate
Impact	Cost: Very Low; Technical Scope: Moderate
Risk Rating	Moderate
Primary Management Strategies	Make salaries as competitive as feasible. Identify promotion opportunities. Develop flexible work schedules. Implement flexibility in work assignments. Training of new staff. Identify staff backups.
Triggers	Staff resignations. Staff reports of receiving outside offers.

Description

This is a period of necessary growth for ALCF as it continues to staff up to operate Mira and Theta together and to advance the ALCF-3 project. An aggressive staff ramp-up, originally planned for FY 2010 through FY 2012, was extended because of budget reductions. An ALCF risk evaluation identified two key risks associated with this ramp-up, and both occurred in CY 2016 as a result of industry competition for retention of existing employees and potential new hires. These two risks have been combined for this discussion, as they are related:

- 25: Challenges encountered in recruiting and hiring new qualified HPC staff.
- 1049: Unable to retain staff due to increased demand for staff with compute expertise and staff worries about DOE funding.

Evaluation

More industry jobs continue to open up that may be attractive to ALCF staff, and positions elsewhere at Argonne also become available. In the past year, one ALCF staff member left, and two moved to positions in the CELS Directorate. Five new full-time staff, one contractor, and two staff transfers (one full-time, one part-time) from other Argonne divisions were added during CY 2016, for a net gain of +3.5 ALCF staff members for the year. A key ALCF leadership position, the Deputy Division Director, was filled from within ALCF during CY 2016. Thus, while ALCF has continued to make good progress on adding new hires, staff retention remains a concern.

Management

Because of industry competition for potential new hires, a limited pool of experienced and available HPC staff, and the fact that candidates do not come out of universities trained for HPC work, it can be very challenging to hire experienced HPC staff. For these reasons, several years ago the ALCF risk management team began preparing to execute mitigations in advance of the occurrence of these risks. When the risks occurred, ALCF was able to continue supporting existing projects successfully even while understaffed.

ALCF has continued to use mitigations to manage both risks over the past year. Facility management continues to replan work as needed, sometimes delaying both planned improvements and lower-priority work. Other mitigation strategies that have been used to address staffing issues include retasking staff, dropping lower-priority tasks, and, when possible, matrixing in staff expertise from other divisions.

By carefully and judiciously managing both risks, ALCF has successfully operated the facility and moved ahead with the ALCF-3 project and with Theta. However, open positions are often difficult to fill, despite aggressive efforts to find and attract qualified candidates, and there continues to be high demand for the skills of ALCF staff members. Thus, both staff recruitment and staff retention will remain a focus for ALCF.

5.3.3 Interruptions to Facility That Provides Cooling

30: Interruptions to Facility That Provides Cooling	
Risk Owner	Mark Fahey
Probability	Very Low
Impact	Cost: Very Low; Technical Scope: Very Low
Risk Rating	Very Low
Primary Management Strategies	Increase redundancy. Site cooling piping is now interconnected to enable chiller plants to provide backup to each other.
Triggers	Rising temperatures of equipment in the Data Center; planned maintenance; monitoring notification of outage or rising temperatures on equipment.

Description

Data Center cooling provided and maintained by Argonne's Facilities Management and Services (FMS) Division and Theory and Computing Sciences (TCS) Building Management was lost on June 6, 2016, when testing by the mechanical contractor of an upgrade to the water pumps caused all Mira cabinets to scram. For further details, see Item 2 in Section 2.1.1.

Evaluation

Several pieces of hardware were damaged as a result of the event; diagnosis to identify all damaged parts took more than week to complete and required three troubleshooting reservations that each took half of Mira out of service for several hours. See further discussion in Item 2 in Section 2.1.1.

Management

An e-mail was sent out to all users explaining the event, and ASCR was notified**.** After being notified that the chiller plant was back functioning normally, ALCF staff waited about 30 minutes to ensure that cooling was stable and then began the process of bringing equipment back online. Once all equipment was back in operation, another e-mail was sent out to users, informing them that service to Mira had been restored. ASCR was notified that ALCF operations had returned to normal.

5.3.4 Changes in Laboratory Indirect Expense (IE) Cost Model, or Increase in IE Rates, Result in an Increased Cost to the ALCF

1090: Changes in Laboratory Indirect Expense (IE) Cost Model, or Increase in IE Rates, Result in an Increased Cost to the ALCF	
Risk Owners	Michael Papka, Darin Wills
Probability	Very Low
Impact	Cost: Moderate; Technical Scope: Very Low
Risk Rating	Low
Primary Management Strategies	Engage laboratory management. Track carefully with DOE and laboratory management. Ensure adequate management reserves. Develop increased austerity plan. Work with laboratory management and DOE to request exemptions and assessment modifications. Monitor indirect expense rate recovery.
Triggers	Monthly indirect expense rate recovery report indicates shortfall. Advance notice from CELS office of possible indirect rate increase. Official notification by laboratory management of an increase in the indirect rates. Laboratory funding is decreasing without a corresponding decrease in laboratory indirect costs.

Description

A change in the division overhead related to the CELS Shared Services model is expected to result in an increase in overhead rates for ALCF that will result in a shortfall of more than $1M in the ALCF FY 2017 budget.

Evaluation

For the FY 2017 budget, carryover funds and contingency funds will cover the shortfall. For out years, a prioritization of spending will need to be completed, and some planned equipment purchases will not be made unless exemptions and/or modified rates are granted or the budget is increased above the currently planned level.

Management

ALCF will track this situation carefully with DOE and with Argonne laboratory management. The division will work with CELS and laboratory management to request exemptions and assessment modifications in order to minimize the impact of this unforeseen change in division overhead.

5.3.5 Facility Power Interruptions

0031: Facility Power Interruptions	
Risk Owner	Mark Fahey
Probability	Low
Impact	Cost: Low; Technical Scope: Low
Risk Rating	Low
Primary Management Strategies	Participation in Data Center management group by the ALCF Director of Operations. ALCF pays part of the cost of an Argonne Data Center liaison. Improve power system bus transfer by investigating the practicality of modifications to the power system to weather bus transfers without interruption.
Triggers	Electrical failure; multiple power quality events; scheduled power outages.

Description

On July 28, 2016, Argonne suffered a partial site-wide power outage that included Building 240. Mira was restored to service after a little more than 28 hours. See further details in Item 4, Section 2.1.1.

Evaluation

There was no loss of data and only minimal damage to Mira hardware. All damaged hardware was replaced during the restoration of the machine.

Management

ALCF began bringing back hardware and services the morning of July 29, 2016, after the cooling and heat load had been brought back to Building 240. Each individual Mira cabinet had to be brought back up separately because of the severity of the power crash.

5.3.6 Interruptions to the Facility Network

0032: Interruptions to the Facility Network	
Risk Owner	William Allcock
Probability	Very Low
Impact	Cost: Very Low; Technical Scope: Very Low
Risk Rating	Very Low
Primary Management Strategies	Computer and Information Services Division will coordinate any planned maintenance or modifications.
Triggers	Planned maintenance; failure of all of the redundant paths.

Description

During a scheduled maintenance period in October 2016, a repair being performed on a Power Distribution Unit that fed the ALCF core router resulted in a false signal to the power controller, indicating that two other power supplies were offline. The controller therefore began turning off units to reduce power consumption, which resulted in a loss of external network connectivity.

Evaluation

Because this event occurred during a scheduled maintenance outage, there were no impacts on users or production. There was no malfunction of the eight router power supplies, and the router needs only four power supplies to run, so without the firmware error, this work would not have resulted in an outage.

Management

The event occurred because of an error in the firmware that provides input to the power controller. After identifying the problem and turning the active power supplies back on, ALCF opened a bug report with the firmware vendor.

5.4 Retired Risks

Two risks were retired during the past year, both as a result of decisions to combine each of them with another risk. These are summarized in Table 5.2.

Table 5.2 Risks Retired During CY 2016

ID	Title	Rating	Management Strategies	Notes
1083	InfiniBand hardware does not meet performance targets.	Low	Accept the lower InfiniBand performance based on application performance. Purchase additional hardware. Work with vendor to tune performance, correct operational issues.	This risk was determined to be covered by Risk 1081, "Storage network hardware is not functional/stable or has performance problems," so Risk 1083 was retired.
1099	INCITE and ALCC do not use all allocated core-hours.	Low	Clear communications between catalysts and projects on importance of time. Ongoing monitoring by catalysts. Adjusting of scheduler priorities to encourage project workflows. Catalysts work with project PI's on any problems with running codes, monitor project usage of allocations. Activate swat team to work with highest priority projects.	This risk was determined to be covered by Risk 1044, "Unable to meet INCITE OMB Metric," so Risk 1099 was retired.

5.5 New and Recharacterized Risks since the Last Review

Staff operating within the ALCF risk culture regularly identify new risks and recharacterize existing risks. This year, the comprehensive review to evaluate all steady-state risks collectively resulted in a larger-than-usual number of risk recharacterizations, which were attributable to scoring adjustments, changes in risk scopes, and consolidation of several risks. In total, no new risks were added but 15 risks were recharacterized in CY 2016. Table 5.3 describes these risks.

Table 5.3 Recharacterized Risks from CY 2016

ID	Title	Rating	Management Strategies	Notes
0025	Staff recruitment challenges	High	Evaluate possible additional recruiting avenues. Prioritize staffing needs. Adjust work planning. Retrain staff to meet ALCF needs. Retask staff as needed. Leave job postings active and open.	A review of historical information concluded that the post-mitigated risk impact scores needed to be raised, resulting in a higher overall score.
1044	Unable to meet INCITE OMB metric	Low	Continuously monitor usage and inform the PIs of their status and impact of project status via bimonthly reports and catalyst weekly reports. Encourage projects to use time earlier in the allocation period. Enable scheduler changes to improve throughput. Include readiness with award. Activate SWAT team to work with highest-priority projects.	The risk scope was expanded to include both the possibility of usage rates being skewed too heavily to the end of the allocation year and of project time requests falling too low.
1050	ALCF has insufficient disk space to support science needs	Low	Storage upgrade. Established and enforced quotas. Require users to more aggressively move data to tape. Pre-negotiate space. Install additional disk. Review allocation proposals for storage needs and evaluate existing storage for space. GPFS HPSS Interface, which converts archive to hierarchal storage.	A review of historical information concluded that the post-mitigated risk scores needed to be lowered, resulting in a lower overall score.
1051	A failure of one minor component can lead to a cascade of failures that becomes a major outage	Very Low	Periodically evaluate systems for potential areas of concern. Develop recovery plans for critical systems. Minimize single points of failure.	A review of historical information concluded that the post-mitigated risk probability score needed to be lowered, resulting in a lower overall score.
1053	Lack of storage systems space (non-HPC)	Very Low	Use a different class of storage as an interim measure. Restrict use of the filers. Purchase additional storage specifically for the filer(s). Determine threshold for concern and put notification methods in place. Regularly review storage usage for data that may be removed.	A review of historical information concluded that the post-mitigated risk probability score and cost score needed to be lowered, resulting in a lower overall score.
1054	Catastrophic failure of home file system	High	HPSS performs automatic aggregation. Tar backup files prior to archiving them. Mirror the data in the home file system. Replicate metadata. Improve performance of the tape system with purchase of additional hardware. Scripts in place to tar the backup files prior to archiving them.	This is a risk that has a low probability but a very high technical impact if it were to occur. The technical impact score was raised to reflect this, resulting in a higher overall score.

Table 5.3 New and Recharacterized Risks from CY 2016 (Cont.)

ID	Title	Rating	Management Strategies	Notes
1056	System stability issues due to upgrades	Low	Perform upgrades on noncritical systems first when feasible. Have a rollback plan in place. Monitor performance closely following upgrade. Work with the vendor to understand the upgrade(s) and the quality control processes. Deep test on test and development systems.	A review of historical information concluded that the post-mitigated risk probability score needed to be lowered, resulting in a lower overall score.
1065	Problems with water cooling	Very Low	Work with vendor to ensure that equipment is appropriately calibrated and monitored. Meet with TCS Building management to help ensure that the management staff understands the system operations. Maintain a second process loop for redundancy. Performing monitoring and deliver notifications of possible issues.	Experience with the water cooling indicates that the probability and technical impacts should all be scored Very Low, resulting in a lowered overall score.
1068	Rack coolant monitor fails to function	Very Low	Consider increasing coolant flow. Contact other users for lessons learned.	Experience with the monitors indicates that the probability and technical impacts should all be scored Very Low, resulting in a lowered overall score
1076	If the Interim Supercomputing Support Facility (ISSF) is decommissioned, ALCF will not have an appropriate facility to host disaster recovery resources	Low	Explore alternative locations onsite. Explore alternative locations offsite (e.g., another lab to host ALCF disaster recovery resources). Accept the risk of a major loss of service or data. Raise concerns with Argonne upper management. Develop plan to move equipment. (At present, there is no planned shutdown of the ISSF.)	If encountered, this risk is easily managed by finding an alternative site, so the technical impact score was reduced to Very Low, resulting in a lowered overall score.
1078	System stability issues	Very Low	Conduct regression testing to monitor system health. Monitor and analyze system failure data.	A review of historical information concluded that the risk probability score needed to be lowered, resulting in a lower overall score.
1079	Compute performance issues	Very Low	Conduct regression testing to monitor system health. Monitor and analyze system failure data. Evaluate initially on test system 2 weeks prior to production.	A review of historical information concluded that the post-mitigated risk impact scores needed to be reduced, resulting in a lower overall score.
1081	Storage network hardware is not functional/stable or has performance problems	Very Low	Test on ALCF test and development platform when feasible. Create/maintain a parts depot on site. Work with vendor.	The risk scope was expanded to cover all storage network hardware, not just InfiniBand.

Table 5.3 New and Recharacterized Risks from CY 2016 (Cont.)

ID	Title	Rating	Management Strategies	Notes
1085	Diagnostic suite and utilities fail to detect hardware problems	Very Low	Track and monitor job and hardware failures, correlate. Work with vendor to resolve issues and improve diagnostic suite.	A review of historical information concluded that the post-mitigated risk probability and technical impact scores should be Very Low, resulting in a lower overall score.
1091	Injury to workers/overall safety of the division	Low	Promote safety culture at all levels of the division. Follow the Argonne Integrated Safety Management (ISM) Plan. Monitor work areas for potential safety concerns. Enforce use of personal protective equipment. Enforce limited access. Redundancy in equipment training.	A review of historical information concluded that the post-mitigated risk probability score should be Very Low, resulting in a lower overall score.

5.6 Projected Major Operating Risks for the Next Year

Table 5.4 lists the current top operating risks projected for CY 2017, along with the current risk rating and management strategies for the risk. These are the risks that experience has shown are most likely to be encountered in any fiscal year.

Table 5.4 Projected Operating Risks for CY 2017

ID	Title	Rating	Management Strategies
1059	Funding/Budget Shortfalls	High	Develop austerity measures. Work closely with DOE sponsors to manage expectations and scope. Plan carefully, in conjunction with program office, for handling Continuing Resolution, leasing costs, and hires. Forward-pay lease to reduce overall leasing costs.
25	Staff Recruitment Challenges	High	Evaluate possible additional recruiting avenues. Prioritize staffing needs. Adjust work planning. Retrain staff to meet ALCF needs. Retask staff as needed.
1049	Staff Retention	Mod	Make salaries as competitive as feasible. Identify promotion opportunities. Develop flexible work schedules. Implement flexibility in work assignments.
1091	Injury to Workers/Overall Safety of the Division	Low	Promote safety culture at all levels of the division. Follow Argonne ISM Plan. Monitor work areas for potential safety concerns. Enforce use of personal protective equipment.

Conclusion

ALCF uses a proven risk management strategy that is documented in its RMP. This document is regularly reviewed and updated to reflect the dynamic nature of risk management, as well as new lessons learned and best practices captured from other facilities. Risk management is a part of the ALCF's culture and applies equally to all staff, from senior management to summer students. A formal risk assessment is performed for every major activity within the ALCF, with informal assessments used for smaller activities. Risks are monitored and tracked using a secure shared cloud-based storage system, plus risk forms and a risk register that are both formatted using Excel. Over the past year, two risks were retired, no new risks were added, and 15 risks were recharacterized. Beyond these activities, many tools are used to manage risks at ALCF, particularly in the area of safety. ALCF's effective risk management plan has contributed to the successful management of all significant risks encountered in the past year.

Section 6. Safety

Has the site implemented measures for safety of staff and the public that are appropriate for HPC/networking facilities?

ALCF Response

The ALCF has an exemplary safety record. Since the division's inception in 2006, the ALCF has never experienced a lost time incident.

The ALCF employs appropriate work planning and control principles. A formal "skill of the worker" document is used for routine tasks. Formal specific procedures are in place for more complex tasks, such as changing out the Blue Gene/Q power supplies (thermal hazard) and node boards (very mild chemical hazard owing to water treatment chemicals, weight, and potential damage to hardware), as well as for performing medium-voltage electrical maintenance. The facility performs hazard analysis and creates work planning and control documents for emergency work or when there is an unexpected change to previously planned work.

This page intentionally left blank.

Section 7. Cyber Security

Have innovations been implemented to improve the facility's cyber security posture? Does the site have a valid cyber security plan and authority to operate?

ALCF Response

Yes. The ALCF improved its cyber security posture. The ALCF worked closely with Argonne's Cyber, Operations, Analysis and Research (COAR) team to navigate the multi-factor authentication (MFA) implementation for Argonne. Level 4 HSPD-12 cards were requested for all privileged users of ALCF. The plan is for privileged staff to begin authenticating to ALCF resources using the HSPD-12 badge starting in January 2017.

There were no cyber security incidents in 2016.

The Argonne Authority to Operate (ATO) includes the ALCF as a major application, and it was granted on November 21, 2016. It is valid as long as Argonne National Laboratory maintains robust, continuous monitoring of the Cyber Security Program as detailed in the letter. A copy of the ATO letter follows.

Department of Energy
Argonne Site Office
9800 South Cass Avenue
Argonne, Illinois 60439

NOV 2 1 2016

Dr. Peter B. Littlewood
Director, Argonne National Laboratory
President, UChicago Argonne, LLC
9700 South Cass Avenue
Argonne, Illinois 60439

Dear Dr. Littlewood:

SUBJECT: AUTHORITY TO OPERATE FOR THE ARGONNE NATIONAL LABORATORY INFORMATION TECHNOLOGY INFRASTRUCTURE

Reference: Letter, J. Livengood to P. Littlewood, dated August 27, 2015, Subject: Renewal of Authority to Operate for the Argonne National Laboratory Information Technology Infrastructure

Over the past year Argonne National Laboratory (Argonne) has modified its Information Technology (IT) Architecture to create a FIPS-199 Low enclave from portions of its previous FIPS-199 Moderate General Computing Enclave. This enclave entitled General Computing – Low contains those portions of the Laboratory IT infrastructure that conduct open science and non-sensitive administrative functions. The remainder of the Laboratory retains its FIPS-199 Moderate rating and is entitled General Computing – Moderate. The only technical distinction between General Computing – Low and General Computing – Moderate is the requirement to employ two-factor authentication within General Computing – Moderate. Thus, no reduction in cyber security was incurred in the creation of General Computing – Low.

The Laboratory has conducted regular continuous monitoring briefings during this re-architecture and has kept me informed of changes in cyber security risk in accordance with the Risk Management Framework. The Laboratory has submitted revised security documentation for the General Computing – Low and General Computing - Moderate enclaves, and has tested seventy seven security controls as part of ANL's 2016 OMB Circular A-123 Compliance Internal Audit. This has demonstrated that the Laboratory's IT Infrastructure is operating at an acceptable level of risk and I am therefore, as the Authorizing Official, approving an Authority to Operate (ATO) for the General Computing – Low enclave and renewing the Authority to Operate for the reconfigured General Computing – Moderate enclave.

The IT Infrastructure continues to contain its sub-component major applications:

- Accelerator Control Systems (APS and ATLAS)
- Argonne Leadership Computing Facility
- Business Systems
- Sensitive Information
- Cyber Federated Model (CFM)

A component of the Office of Science

NOV 2 1 2016

Four of the five major applications have components in both the General Computing – Low and General Computing Moderate enclaves.

This ATO will remain in effect as long as the Laboratory carries out continuous monitoring under the Risk Management Framework and there are no significant changes to Argonne's IT Infrastructure. The Laboratory should retain a copy of this letter with the security authorization package. If I can be of any assistance, please contact me or have your staff contact Francis Healy at (630) 252-2827 or e-mail frank.healy@science.doe.gov.

Sincerely,

Joanna M. Livengood
Manager

cc: S. Hannay, ANL-CIS
M. Skwarek, ANL-CIS
M. Kwiatkowski, ANL-CIS
V. Dattoria, SC-21.2
N. Masincupp, SC-OR
F. Healy, SC-CH

This page intentionally left blank.

Section 8. Summary of the Proposed Metric Values for Future OARs

Are the performance metrics used for the review year and proposed for future years sufficient and reasonable for assessing Operational performance?

ALCF Response

The ALCF and the DOE have agreed to the 2017 metrics and targets as proposed in the March 2016 OAR report. The proposed metrics and targets are reasonable measures of facility performance that are consistent with metrics and targets used at other facilities. For 2018, the proposed metrics and targets for the current production resources remain the same as for 2017. A new system, Theta, is expected to enter production in 2017; therefore, 2017 metrics and targets and 2018 proposed metrics and targets for this system have been added. DOE and the ALCF have agreed that Theta will have no INCITE hours during 2017; therefore, there is no capability metric for Theta in 2017.

8.1 Overview

The ALCF metrics and targets are reasonable measures of facility performance that are consistent with metrics and targets used at other facilities. For 2018, the proposed metrics and targets for the current production resource, Mira, will remain the same as for 2017. Appropriate metrics and targets have been proposed for Theta, a resource expected to enter production in 2017. The 2017 metrics are covered in Section 8.2 and the 2018 metrics are covered in Section 8.3.

8.2 ALCF 2016 OA Performance Metrics

The Operational Assessment performance metrics, 2016 targets and actuals, and agreed-upon 2017 targets are presented in Table 8.1. The new ALCF resource, Theta, is planned to go into production in July 2017 and provide hours for the ALCC allocation program. ASCR has requested that all user facilities use a target of 90 percent for Scheduled Availability for the lifetime of the production resources; therefore, the 2017 Schedule Availability target for Theta will be 90%. DOE and the ALCF have agreed that Theta will provide no INCITE hours during 2017; therefore, there is no capability metric for Theta in 2017.

Table 8.1 Performance Metrics: 2016 Targets, 2016 Actuals, and Agreed-Upon 2017 Targets

Area	Metric	2016 Targets	2016 Actuals	2017 Targets
User Results	User Survey – **Overall Satisfaction**	3.5/5.0	4.5/5.0	3.5/5.0
	User Survey – **User Support**	3.5/5.0	4.5/5.0	3.5/5.0
	User Survey – **Problem Resolution**	3.5/5.0	4.5/5.0	3.5/5.0
	User Survey – **Response Rate**	25.0%	45.1%	25.0%
	% User Problems Addressed Within Three Working Days	80.0%	95.9%	80.0%
Business Results	Mira Overall Availability	90.0%	94.9%	90.0%
	Mira Scheduled Availability	90.0%	98.9%	90.0%
	% of INCITE core-hours from jobs run on 16.7% or more of Mira (131,072–786,432 cores)	40.0%	76.4%	40.0%
	% of INCITE core-hours from jobs run on 33.3% or more of Mira (262,144–786,432 cores)	10.0%	40.9%	10.0%
	Theta Overall Availability (July 2017–December 2017)	N/A	N/A	80.0%
	Theta Scheduled Availability (July 2017–December 2017)	N/A	N/A	90.0%

8.3 ALCF Proposed 2018 OA Performance Metrics

The Operational Assessment performance metrics, agreed-upon 2017 targets, and 2018 proposed targets are shown in Table 8.2. It has not yet been determined whether Theta will provide INCITE hours during 2018. A potential capability metric has been proposed in case there are INCITE allocations; it will be removed in the event that DOE and the ALCF agree that Theta will not be part of the INCITE allocation program in 2018.

Table 8.2 Performance Metrics: Agreed-Upon 2017 Targets and Proposed 2018 Targets

Area	Metric	2017 Targets	Proposed 2018 Targets
User Results	User Survey – **Overall Satisfaction**	3.5/5.0	3.5/5.0
	User Survey – **User Support**	3.5/5.0	3.5/5.0
	User Survey – **Problem Resolution**	3.5/5.0	3.5/5.0
	User Survey – **Response Rate**	25.0%	25.0%
	% User Problems Addressed within Three Working Days	80.0%	80.0%

Table 8.2 Performance Metrics: Agreed-Upon 2017 Targets and Proposed 2018 Targets (Cont.)

Area	Metric	2017 Targets	Proposed 2018 Targets
Business Results	Mira Overall Availability	90.0%	90.0%
	Mira Scheduled Availability	90.0%	90.0%
	% of INCITE core-hours from jobs run on 16.7% or more of Mira (131,072–786,432 cores)	40.0%	40.0%
	% of INCITE core-hours from jobs run on 33.3% or more of Mira (262,144–786,432 cores)	10.0%	10.0%
	Theta Overall Availability	80.0%	90.0%
	Theta Scheduled Availability	90.0%	90.0%
	% of INCITE core-hours from jobs run on 20% or more of Theta (41,472–207,360)	N/A	20.0%

8.4 ALCF Reportable-Only Metrics (No Targets)

ALCF has a set of metrics that have no targets and are only reported. These are shown in Table 8.3.

Table 8.3 ALCF Reportable-Only Metrics

Area	Metric (No Targets)
User Support Results	Summarize training events and provide examples of in-depth collaborations between facility staff and the user community
Business Results	Report MTTI, MTTF, Utilization, and Usage for the past CY
INCITE Management	Report reviewer survey responses and the proposal allocation results (number of proposals, number of awards, percentage awarded, number of hours requested/awarded, oversubscription) to DOE.
Science Results	Track and report the number of publications written annually (projects are tracked for five years after award). Report on at least five significant scientific accomplishments and the DD awards.
Innovation	Report on innovations that have improved operations.

Conclusion

The agreed-upon 2017 metrics and targets are reasonable measures of facility performance that are consistent with metrics and targets used at other facilities. For 2018, the proposed metrics and targets will remain the same as for 2017 for the current production resources. For the future system Theta, which is expected to enter production in July 2017, a set of metrics and targets in line with past new systems have been agreed upon for 2017 and proposed for 2018. Achieving the agreed-upon 2017 and the proposed 2018 targets will indicate that the facility is performing up to stakeholder expectations. ALCF anticipates being able to meet all metric targets.

This page intentionally left blank.

Appendix A – ALCF Director's Discretionary Projects

January 1, 2016 – December 31, 2016

Mira DD

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
3D_MOC_Reactor_Sim	John Tramm	Massachusetts Institute of Technology	Domain Decomposed 3D Method of Characteristics Reactor Simulation Study	Nuclear Energy	1,000,000
Abeta42	Georges Belfort	Rensselaer Polytechnic Institute	Investigating a Drug Target of the Amyloid Beta (Aβ) Peptide for Alzheimer's Disease with Molecular Simulations	Biological Sciences	2,000,000
Acceptance	Ti Leggett	Argonne National Laboratory	Acceptance Tests	Internal	2,000,000
Aeroacoustics-Optics	Meng Wang	University of Notre Dame	Turbulence Simulations for Aeroacoustics and Aero-Optics	Engineering	2,000,000
AIR	Ilpo Vattulainen	University of Helsinki / Tampere University of Technology	Insulin-Induced Activation of the Insulin Receptor	Biological Sciences	100,000
ALCF_Getting_Started	Ray Loy	Argonne National Laboratory	ALCF Getting Started	Training	470,000
ALCF_Scaling_Event	Chel Lancaster	Argonne National Laboratory	ALCF_Scaling_Event	Training	5,100,000
AlgLA	Oded Schwartz	The Hebrew University	Algorithmic Linear Algebra	Computer Science	250,000
Allinea	Ray Loy	Argonne National Laboratory	Improved Debugging Memory Usage for BG/Q	Internal	2,000,000
Alloy-Corrosion	Hendrik Heinz	University of Colorado Boulder	Understanding and Designing Corrosion-Resistant Alloys	Materials Science	4,000,000
alpha-nek	Maxwell Hutchinson	The University of Chicago	DNS of Multi-Mode Rayleigh-Taylor Instability	Engineering	5,048,576
AMOEBA_Swift	Ying Li	Argonne National Laboratory	AMOEBA Potential Parallel Parameterization Using Swift on BG/Q	Biological Sciences	300,000
Angora_scaling_study	Allen Taflove	Northwestern University	Angora Scaling Study	Biological Sciences	6,000,000
APS_UBeam_Dynamics	Michael Borland	Argonne National Laboratory : Advanced Photon Source	Beam Dynamics Simulations for the Advanced Photon Source Upgrade	Physics	22,000,000
ARGO	Pete Beckman	Argonne National Laboratory	ARGO	Computer Science	20,000
AT1R_activation	Xavier Deupi	Paul Scherrer Institute	Study of the Mechanism of Activation of the Angiotensin II Type 1 Receptor Using Enhanced Molecular Dynamics Simulations	Biological Sciences	2,000,000
ATPESC16_Instructors	Marta Garcia	Argonne National Laboratory : Leadership Computing Facility	Argonne Training Program on Extreme Scale Computing for All Instructors	Training	1,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
ATPESC2016	Marta Garcia	Argonne National Laboratory : Leadership Computing Facility	Argonne Training Program on Extreme Scale Computing	Training	10,000,000
aurora_app	William Scullin, Kevin Harms	Argonne National Laboratory	Aurora Application Enablement	Computer Science	1,200,000
Bachalo_Johson_DNS	Philippe R. Spalart	The Boeing Company	Direct Numerical Simulation of Bachalo-Johnson Transonic Separated Flow	Engineering	6,000,000
BGQ_Energy_Profiling	Venkatram Vishwanath	Argonne National Laboratory	Understanding and Characterizing the Power Consumption and Energy Efficiency of Applications on BG/Q	Computer Science	1,000,000
BigDFT4Q	Alvaro Vazquez-Mayagoitia	Argonne National Laboratory	Simulation of Large Molecular Systems with BigDFT	Chemistry	2,000,000
BrainImagingADSP	Doga Gursoy	Argonne National Laboratory	Large-Scale Computing and Visualization on the Connectomes of the Brain	Biological Sciences	10,000,000
Broadband_Fan_Noise	Daniel J. Bodony	University of Illinois at Urbana-Champaign	Sources of Broadband Fan Noise in the NASA/GE Source Diagnostic Test Fan	Engineering	2,000,000
Bubblecollapse	Eric Johnsen	University of Michigan	Dynamics and Heat Transfer in Bubble Collapse near Solid Surfaces	Engineering	1,500,000
Camellia	Nathan Roberts	Argonne National Laboratory : Leadership Computing Facility	Camellia for Discontinuous Petrov-Galerkin Simulations of Incompressible Flow	Physics	5,000,000
Catalyst	Katherine Riley	Argonne National Laboratory	Catalyst	Internal	20,000,000
CatalystInt_DD	Frank Abild-Pedersen	SLAC National Accelerator Laboratory	Catalyst Support Interactions	Chemistry	6,000,000
CEED_ECPAD	Misun Min	Argonne National Laboratory	CEED	Mathematics	1,000,000
CERES	Andrew Chien	The University of Chicago	CERES	Computer Science	2,500,000
CESM_Atmos_PostProc	Warren Washington	The National Center for Atmospheric Research	CESM_Atmos	Earth Science	10,500,000
CharmRTS	Laxmikant V. Kale	University of Illinois at Urbana-Champaign	Charm++ and its Applications	Computer Science	2,000,000
cmsframemini	Elizabeth Sexton-Kennedy	Fermilab	CMS Framework MiniApp	Physics	10,000
CMSHPCGrid	Harvey Newman	California Institute of Technology	Integrating HPC to CMS Production	Physics	400,000
CMSHPCProd	Harvey Newman	California Institute of Technology	CMS Production on HPC	Physics	800,000
CMSHPCSherpack	Harvey Newman	California Institute of Technology	Sherpa for CMS	Physics	400,000
CMT	Scott Parker	Argonne National Laboratory	Compressible Multiphase Turbulence	Engineering	3,000,000
CNTmetallization	Iman Salehinia, Mike Papka	Northern Illinois University	Metallization of CNT for Thermal and Structural Applications	Materials Science	500,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CompBIO	Rick Stevens	Argonne National Laboratory / The University of Chicago	Multiscale Simulations in Biology: Evolution and Ecology of Microbes, Cancer Research, Machine Learning, Infectious Disease	Biological Sciences	1,000,000
ComplexTurbulence	Krishnan Mahesh	University of Minnesota	High-fidelity Simulations of Complex Turbulent Flows	Engineering	7,000,000
Compressible_RT	Hussein Aluie	University of Rochester	Multi-Scale Coupling in Rayleigh-Taylor Flows	Engineering	700,000
CORALDev	Scott Parker	Argonne National Laboratory	CORAL Development and Testing	Internal	2,000,000
CORALtestApps	James Osborn	Argonne National Laboratory	Preparing Test Applications for CORAL Machines	Internal	200,000
Cray	Ti Leggett, Mark Fahey, Susan Coghlan	Cray Inc.	Cray Installation	Internal	1,000,000
critical_perf	Kalyan Kumaran, Ray Loy	Argonne National Laboratory	Critical Debugging Project	Internal	80,000,000
datascience	Venkatram Vishwanath	Argonne National Laboratory	ALCF Data Science and Workflows Allocation	Internal	2,000,000
DetailedKinetics	William Anderson	Purdue University	Combustion Instability Simulations with Detailed Kinetics for Direct Comparison with Experiments	Chemistry	2,000,000
DFT_scaling_Mira	S. Pamir Alpay	University of Connecticut	Quantum Mechanical Understanding of Complex Correlated Materials	Materials Science	1,000,000
DiscoveryEngines	Justin M Wozniak, Rajkumar Kettimuthu	Argonne National Laboratory	Integrating Simulation and Observation: Discovery Engines for Big Data	Materials Science	10,100,000
DNS-ABF	Antonino Ferrante	University of Washington	Direct Numerical Simulation of an Aft-Body Flow	Engineering	1,000,000
DNSGeo	Christos Frouzakis	University of Western Macedonia, Greece	DNS of Forced- and Auto-Ignition in Spherical and Engine-like Geometries	Chemistry	450,000
DSEM_SUPERSONIC	Farzad Mashayek	University of Illinois at Chicago	Simulation of Supersonic Combustion	Engineering	2,000,000
duanl	Lian Duan	Missouri University of Science and Technology	Numerical Simulation of Acoustic Radiation from High-Speed Turbulent Boundary Layers	Engineering	3,000,000
dvdt	William E. Allcock	Argonne National Laboratory	Accelerating the Rate of Progress towards Extreme Scale Collaborative Science	Computer Science	5,000,000
EarlyPerf_theta	Kalyan Kumaran	Argonne National Laboratory : Leadership Computing Facility	Enabling Science on Theta: Tools Libraries Programming Models & Other System Software	Internal	1,000,000
EKaos	Ali Mani	Stanford University	DNS of Chaotic Electroconvection	Engineering	1,000,000
em-brain	Mike Papka	Argonne National Laboratory	ANL-BrainProject	Biological Sciences	20,000
es_tddft	Yosuke Kanai	The University of North Carolina at Chapel Hill	First-Principles Simulation of Electronic Excitation Dynamics in Liquid Water and DNA under Proton Irradiation	Chemistry	3,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
ExaHDF5	Venkatram Vishwanath	Argonne National Laboratory	ExaHDF5: Advancing HDF5 HPC I/O to Enable Scientific Discovery	Computer Science	6,262,092
EXCEL	Francisco Doblas-Reyes	Barcelona Supercomputing Center	EXtreme Climate Event attribution using dynamical seasonaL predictions	Earth Science	300,000
ExM	Justin M Wozniak	Argonne National Laboratory : Mathematics and Computer Science	Extreme Many-Task Computing with Swift	Computer Science	101,000
Extreme_Scale_TS	William M. Tang	Princeton Plasma Physics Laboratory	Extreme Scale Turbulence Simulations	Physics	15,000,000
FokkerPlanck	Debojyoti Ghosh	Argonne National Laboratory	Fokker Plank Model of a 3-Bus Power System	Energy Technologies	2,621,440
folding	Sichun Yang	Case Western Reserve University	Folding and Stability of an Intrinsically Disordered Domain in Estrogen Receptor	Biological Sciences	5,000,000
FPMC	Neeraj Rai	Mississippi State University	First Principle Monte Carlo Algorithm Development and Implementation in CP2K	Chemistry	5,000,000
Framework_dynamics	Karl Andrew Wilkinson	University of Cape Town	Dynamic Properties of Porous Frameworks upon the Absorption of Gas Molecules	Chemistry	8,000,000
FTQCD	Xiao-Yong Jin	Argonne National Laboratory	Finite Temperature Lattice QCD Startup	Physics	2,000,000
GalaxiesOnFire	Philip F. Hopkins	California Institute of Technology	Galaxies on FIRE: Shedding Light on Dark Matter	Physics	2,000,000
GAtor	Noa Marom	Carnegie Mellon University	GAtor: A Cascade Genetic Algorithm for Crystal Structure Prediction	Materials Science	5,000,000
GenomeOrganization	Alexey V. Onufriev	Virginia Polytechnic Institute and State University	Investigation of Genome Compaction and Organization with All-Atom and Coarse-Grained MD Simulation	Biological Sciences	900,000
graph500	Andrew Lumsdaine	Indiana University	Graph500 Benchmark Run on Intrepid	Computer Science	750,000
GRChombo	Hal Finkel	Argonne National Laboratory	Early Universe Phase Transitions with Strong Gravity and Instabilities in Higher-Dimensional Black Holes	Physics	2,000,000
Hephaestus	Andrey Beresnyak	Naval Research Laboratory	Testing of New PIC-MHD Code Hephaestus	Physics	1,000,000
HighAspectRTI	Maxwell Hutchinson	The University of Chicago	Direct Numerical Simulation of the High-Aspect Rayleigh-Taylor Instability	Engineering	4,600,000
HighReyTurb_PostProc	Robert D. Moser	The University of Texas at Austin	Data Analysis of Turbulent Channel Flow at High Reynolds Number	Engineering	14,500,000
HiPPSTR4096	Andrew D Bragg	Duke University	High-Resolution DNS of Inertial Particle Motion in Turbulence	Engineering	1,800,000
Hoskin_SU2	Dominique Hoskin	Argonne National Laboratory	SU2 Summer Internship	Engineering	1,000,000
HPC-Exp-Cancer	Eric Stahlberg	Frederick National Laboratory for Cancer Research	HPC Explorations Supporting Cancer Research	Biological Sciences	5,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
HPCTuning	Khaled Ibrahim	Lawrence Berkeley National Laboratory	HPC Applications Tuning	Computer Science	10,000,000
hpMusic	ZJ Wang	University of Kansas	A Study of Scalability of hpMusic	Engineering	800,000
HRDBMS	Jason Arnold	Illinois Institute of Technology	HRDBMS: NewSQL Analytics Database	Computer Science	10,000
HSW_TURBULENCE	Vadim Roytershteyn	Space Science Institute	Hybrid Simulations of Solar Wind Turbulence	Physics	1,500,000
Hybrid-C-Modelling	Jinxun Liu	U.S. Geological Survey / Western Geographic Science Center	Simulating Global Terrestrial Carbon Sequestration and Carbon Transport to Aquatic Ecosystems - Pilot Study	Earth Science	2,000,000
HydraHPC	Luigi Capone	Rolls Royce plc	Hydra App Benchmark	Engineering	200,000
IME_BlockCoPolymers	Venkatram Vishwanath	The University of Chicago / Argonne National Laboratory : Institute of Molecular Engineering/ Materials Science	Scalable Data Analysis of Soft X-Ray Scattering for APS Beamline Experiments	Materials Science	100,000
Injfuel	Marco Arienti	Sandia National Laboratories	High-Fidelity Characterization of Fuel Injection	Engineering	300,000
ITHALES	Koen Hillewaert	Cenaero	Industrial Turbomachinery High Accuracy LES – iTHALES	Engineering	2,000,000
JetBlending	Dee Dickerson	The Dow Chemical Company	A Computational Study of Jet Blending in Large Vessels	Engineering	1,000,000
LESOIF_2015	Francesco Grasso	DynFluid Laboratory - Arts et Métiers Paris Tech	LES of Shock-Wave Boundary Layer Interaction in Internal Flow with Corner Effects	Engineering	3,800,000
LES_Cooling	Gustavo Ledezma	General Electric Company Global Research	Large-Eddy Simulation of Turbine Internal Cooling Passages	Engineering	1,000,000
Li-S_Trans	Ying Li	Argonne National Laboratory : Leadership Computing Facility	Ion Transport in Li-S solid	Materials Science	15,000,000
LiquidGaNa_PostProc	Janet Scheel, Joerg Schumacher	Occidental College	Convective Turbulence in Liquid Gallium and Sodium	Engineering	3,000,000
lossynek	Oana Marin	Argonne National Laboratory : Mathematics and Computer Science (MCS)	Lossy Data Compression for Large-Scale CFD simulations	Engineering	2,000,000
Low_Pt_content_catal	Binay Prasai	Central Michigan University	Modeling and Prediction of Nanocatalyst for Fuel Cells	Energy Technologies	2,000,000
LQCDdev	James Osborn	Argonne National Laboratory	Lattice QCD Development	Physics	1,000,000
LTC_Aramco	Sibendu Som	Argonne National Laboratory/Aramco Services Company : Leadership Computing Facility	Investigation of a Low Octane Gasoline Fuel for a Heavy-Duty Diesel Engine in a Low-Temperature Combustion Regime	Engineering	3,000,000
LUCMDproject	Nikolai Smolin	Loyola University Chicago	Molecular Dynamics Simulation of SERCA Troponin Complex and TRIM5a	Biological Sciences	368,640

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
Maintenance	Mark Fahey	Argonne National Laboratory	LCF Operations System Maintenance	Internal	212,000,000
MBPolOpt	Francesco Paesani	University of California San Diego	Enabling Chemical Accuracy through Many-Body Molecular Dynamics	Chemistry	1,000,000
MCMD_MSM	Phay Ho	Argonne National Laboratory	Efficient Computational Algorithms for Modeling Intense X-Ray Interaction with Nanosized Systems	Physics	1,000,000
Meso_CCS_DD16	Roberto Paoli	CERFACS	Contrail-Cirrus Sensitivity to Turbulence Fluctuations using LES Data from Meso_CCS_13 INCITE Project	Earth Science	1,000,000
Metal-sulfur_protein	Sergey Varganov	University of Nevada, Reno	Modeling Nonadiabatic Spin-Forbidden Reaction Mechanisms in Metal-Sulfur Proteins	Chemistry	3,000,000
MFPM	Olivier Desjardins	Cornell University	Large-scale Simulation of Multiphase Flow in Porous Media	Engineering	1,000,000
MHD-turb	Pui-kuen Yeung	Georgia Institute of Technology	Strained and Magnetohydrodynamic Turbulence	Engineering	1,000,000
MICCoM	Justin Wozniak	Argonne National Laboratory : Mathematics and Computer Science	Midwest Integrated Center for Computational Materials	Materials Science	50,000
MLGAOpt	Paul C. Jennings	SLAC National Accelerator Laboratory	Optimization and Benchmarking of a Machine Learning Genetic Algorithm	Chemistry	1,000,000
MM-MEDE	Mauricio Ponga	The University of British Columbia	Multiscale Modeling of Materials under Extreme Dynamic Environments through Large-scale Computer Simulations	Materials Science	2,000,000
MonteCarloSMR	Paul Romano	Argonne National Laboratory	Monte Carlo Particle Transport for Small Modular Reactors	Nuclear Energy	2,000,000
MPAccelMod	James F. Amundson	Fermilab	Massively Parallel Accelerator Modeling	Physics	1,000,000
mpich-bgq	William Scullin	Argonne National Laboratory	MPICH Testing on BG/Q	Computer Science	300,000
MPICH_MCS	Ken Raffenetti	Argonne National Laboratory : Mathematics and Computer Science	MPICH - A High Performance and Widely Portable MPI Implementation	Computer Science	10,000,000
NAMD_FEP_validation	Prabhu Raman	Dassault Systemes BIOVIA	High Throughput Relative Free Energy Computations	Biological Sciences	5,000,000
nci_doe_pilot1	Rick Stevens	Argonne National Laboratory	nci_doe_pilot1	Biological Sciences	100,000
NekPerformance	Oana Marin	Argonne National Laboratory : Mathematics and Computer Science	Nek Performance Evaluation	Computer Science	1,000,000
NRCM_DD	V. Rao Kotamarthi	Argonne National Laboratory	Dynamic Downscaling of Climate Models	Earth Science	1,000,000
NUMA	Andreas Mueller	Naval Postgraduate School	Scalability Study for NUMA (Non-hydrostatic Unified Model of the Atmosphere)	Earth Science	1,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
Omega-NIF_Exp	Don Q. Lamb	The University of Chicago / Argonne National Laboratory	Mira Simulations of High-Intensity Laser Experiments to Study Turbulent Amplification of Magnetic Fields	Physics	10,000
OpenFOAM-ALCF	Ramesh Balakrishnan	Argonne National Laboratory	OpenFOAM-based Computational Fluid Dynamics Simulations at the Argonne Leadership Computing Facility	Engineering	5,000,000
Operations	Mark Fahey	Argonne National Laboratory	Systems Administration Tasks	Internal	20,000,000
Order-disorder_2D	Salvador Barraza-Lopez	University of Arkansas	Order-Disorder Transitions in Two-dimensional Atomic Materials and Effects on Materials Properties	Physics	2,000,000
OSCon	Andreas Glatz	Argonne National Laboratory	Optimizing Superconductor Performance through Large-Scale Simulation	Materials Science	250,000
P3DFFT	Dmitry Pekurovsky	University of California San Diego	Performance Studies of Three-dimensional Fast Fourier Transforms using Overlap of Communication with Computation	Computer Science	1,500,000
ParBous_ProcTrans	Susan Kurien	Los Alamos National Laboratory	New Post-Processing Diagnostics of Boussinesq Flow Data	Engineering	500,001
Particle_Flow	Brian Helenbrook	Clarkson University	Examination of Particle-Wall Collision Models in Turbulent Particle Laden Flows	Engineering	500,000
PBSim	Lambert Fick	Texas A&M University	DNS of Incompressible Flows in Packed Pebble Beds	Engineering	7,000,000
Performance	Kalyan Kumaran, Ray Loy	Argonne National Laboratory	Performance	Internal	40,000,000
PHASTA_NCSU	Igor A. Bolotnov	North Carolina State University (NCSU)	Multiphase Simulations of Nuclear Reactor Thermal Hydraulics	Engineering	13,000,000
PICSSAR	Jean-Luc Vay	Lawrence Berkeley National Laboratory	Particle-In-Cell Scalable Spectral Relativistic	Physics	5,250,000
pl_binding_with_fmo	Casper Steinmann Svendsen	The University of Bristol	Scaling of the FMO Method for Heterogeneous Systems	Chemistry	2,000,000
Polymer_Deformation	Mark Horstemeyer	Mississippi State University	Reactive Molecular Dynamics Simulations of Polyethylene Deformation	Materials Science	2,000,000
PROTEUS_2016	Emily Shemon	Argonne National Laboratory	NEAMS Neutronics Verification and Validation Simulations	Nuclear Energy	13,177,588
PTPases	Andrei Karginov	University of Illinois at Chicago	Study of Protein Tyrosine Phosphatases	Biological Sciences	1,000,000
Py8HPC	Stephen Mrenna	Fermilab	Py8HPC	Physics	100,000
QCDHPC	Radja Boughezal	Argonne National Laboratory : High Energy Physics	Predicting the Terascale On-Demand with High Performance Computing	Physics	9,000,000
QETSc	Murat Keceli	Argonne National Laboratory	Spectrum Slicing Eigensolver for Ab Initio Simulations	Chemistry	750,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
QMC-Alloying	Jeffrey C Grossman	Massachusetts Institute of Technology	Alloying Li Based High-Capacity Cathode Materials using Quantum Monte Carlo	Materials Science	10,000,000
QMC-Prep	Anouar Benali	Argonne National Laboratory	QMC Umbrella	Chemistry	10,000,000
QMCMag	Lucas K. Wagner	University of Illinois at Urbana-Champaign	Using Quantum Monte Carlo for Magnetic Materials	Physics	7,000,000
qmcspinvdwsolid	Lubos Mitas	North Carolina State University	Quantum Monte Carlo for Spin-Orbit Interactions Spintronic and van der Waals systems	Physics	2,000,000
QuanPol	Hui Li	University of Nebraska-Lincoln	QuanPol QMMM Style MP2 Simulation Methods	Chemistry	2,000,000
QuantumDS	Alvaro Vazquez-Mayagoitia	Argonne National Laboratory	Quantum Mechanics and Data Science	Chemistry	2,000,000
Quinoa	Jozsef Bakosi	Los Alamos National Laboratory	Asynchronous Navier-Stokes Solver on 3D Unstructured Grids for the Exascale Era	Engineering	2,000,000
radix-io	Philip Carns	Argonne National Laboratory	System Software to Enable Data-Intensive Science	Computer Science	2,000,000
Rayleigh-training	Nicholas Featherstone	University of California Davis	An Introduction to Rayleigh Using Mira	Earth Science	1,000,000
rcduct	Hassan M. Nagib	Illinois Institute of Technology	Influence on Duct Corner Geometry on Secondary Flow: Convergence from Duct to Pipe Flow	Engineering	4,600,000
ReactingRT	Praveen Ramaprabhu	University of North Carolina at Charlotte	New Pathways to Stability and Instability in Rayleigh-Taylor Non-Premixed Flames	Engineering	2,000,000
ReaxCath	Christopher Knight	Argonne National Laboratory	Reactive Modeling of Battery Cathodes and Interfaces	Chemistry	2,000,000
rec_sironi	Lorenzo Sironi	Columbia University	Particle-in-Cell Simulations of Explosive Reconnection in Relativistic Magnetically-Dominated Plasmas	Physics	1,000,000
REEs_and_actinides	Deborah A. Penchoff	Institute for Nuclear Security	Accelerating Selective Binding of Rare Earth Elements and Actinides	Chemistry	2,000,000
ReliableEnergySys	Cosmin G. Petra	Argonne National Laboratory : Mathematics and Computer Science	Assessing Reliability of Converging Energy Infrastructure Systems	Mathematics	4,000,000
RESTMD_DD14	Tom Keyes	Boston University / Louisiana State University	Replica Exchange Statistical Temperature Molecular Dynamics in LAMMPS	Chemistry	4,000,000
RNA-stcalculation	Yun-Xing Wang	National Cancer Institute / National Institutes of Health	Computing Three-Dimensional Structures of Large RNA from Small Angle X-Ray Scattering Data and Secondary Structures	Biological Sciences	25,000,000
rtflames	Elizabeth P. Hicks	Epsilon Delta Labs	DNS Simulations of Turbulent Rayleigh-Taylor Unstable Flames using Nek5000	Physics	800,000
SciDAC_PILOT_Fission	Brian Wirth	University of Tennessee, Knoxville	Advancing Understanding of Fission Gas Behavior in Nuclear Fuel Through Leadership Class Computing	Nuclear Energy	1,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
SC_Architectures_DD	Tjerk Straatsma, Katerina Antypas, Tim Williams	Oak Ridge National Laboratory	Portable Application Development for Next Generation Supercomputer Architectures	Computer Science	10,000,000
SDAV	Mike Papka, Joe Insley	Argonne National Laboratory : Computing, Environment and Life Sciences	SciDAC Scalable Data Management Analysis and Visualization	Computer Science	3,000,000
SENSEI	Venkat Vishwanath	Argonne National Laboratory	Scalable Analysis Methods and In Situ Infrastructure for Extreme Scale Knowledge Discovery	Computer Science	3,000,000
Shell_Cat	Leonardo Spanu	Shell International E&P, Inc.	Investigation of Catalytic Properties of Nanoclusters	Chemistry	2,000,000
SherpaProd	Taylor Childers	Argonne National Laboratory	Sherpa Parton Generator Production	Physics	5,000,000
SolarWindowsADSP	Jacqueline Cole	University of Cambridge	Data-Driven Molecular Engineering of Solar-Powered Windows	Materials Science	110,000,000
Sparse-Solvers	Luke Olson	University of Illinois at Urbana-Champaign	Reducing Communication in Sparse Solvers	Computer Science	200,000
SSSPScalability	David Padua	University of Illinois at Urbana-Champaign	Iterative Stepping: A Faster Algorithm and Implementation for Single-Source Shortest Path and Betweenness Centrality	Computer Science	264,081
Thermal_Transport	Lucas Lindsay	Oak Ridge National Laboratory	First Principles Thermal Transport for Basic Energy Sciences	Physics	500,000
Tools	Scott Parker	Argonne National Laboratory : Leadership Computing Facility	ALCF Performance Tools	Internal	5,000,000
TopologyMapping	Zhiling Lan	Illinois Institute of Technology	Topology Mapping of Irregular Applications	Physics	1,619,200
TotalView	Peter Thompson, Ray Loy	Rogue Wave Software, Inc.	TotalView Debugger on Blue Gene P	Internal	250,000
TurbulentLiquidDrop	Arne J. Pearlstein	University of Illinois at Urbana-Champaign	Computation of Transitional and Turbulent Drop Flows for Liquid Carbon Dioxide Drops Rising in Seawater	Engineering	1,000,000
umncafi	Joseph W. Nichols	University of Minnesota	Large Eddy Simulation for the Prediction and Control of Impinging Jet Noise	Engineering	5,000,000
UTRC-Turbine	Solkeun Jee	United Technologies Research Center Inc.	High-Fidelity Simulation of Turbines in Engine-Relevant Conditions Toward Next-Generation Turbine Designs	Engineering	1,000,000
VarRhoFlow	Paul E. Dimotakis	California Institute of Technology	Variable-Density Fluid Dynamics	Engineering	3,000,000
VDVAT	Daniel Livescu	Los Alamos National Laboratory	Variable-Density Turbulence under Variable Acceleration	Engineering	1,500,000
Vendor_Support	William E. Allcock	Argonne National Laboratory	Vendor Support	Internal	1,000,000
VERIFI_Workshop	Sibendu Som	Argonne National Laboratory	VERIFI Workshop for High-Performance Computing-Enabled Engine Simulations	Engineering	2,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
VirtualEye	Marco L. Bittencourt	University of Campinas (UNICAMP)	Computational Modeling of the Human Eye	Engineering	7,000,000
visualization	Mike Papka	Argonne National Laboratory	Visualization and Analysis Research and Development for ALCF	Internal	8,639,580
Viz_Support	Joe Insley	Argonne National Laboratory	Visualization Support	Computer Science	100,000
VSFi-DNS	Rafael Tinoco	University of Illinois at Urbana-Champaign	Vegetation-Sediment-Flow Interactions: DNS of Turbulent Oscillatory Flow and Sediment Transport on Aquatic Ecosystems	Engineering	200,000
WaterHammer	Hong Zhang	Argonne National Laboratory : Mathematics and Computer Science	Water Hammer Simulation	Mathematics	30,000
wdmerger	Maximilian Katz	Stony Brook University	White Dwarf Mergers on Adaptive Meshes	Physics	2,000,000
WOInterface	Baofu Qiao	Argonne National Laboratory	Interfacial Behavior of Alcohol at Water/Organic Biphasic System	Computer Science	5,000,000
				Total Mira DD	**1,026,672,198**