ANL/ALCF-20/04

2017 Operational Assessment Report

ARGONNE LEADERSHIP COMPUTING FACILITY

On the cover: This visualization shows a small portion of a cosmology simulation tracing structure formation in the universe. The gold-colored clumps are high-density regions of dark matter that host galaxies in the real universe. This BorgCube simulation was carried out on Argonne National Laboratory's Theta supercomputer as part of the Theta Early Science Program.

Image: Joseph A. Insley, Silvio Rizzi, and the HACC team, Argonne National Laboratory

Contents

Contents (Cont.)

Contents (Cont.)

Contents (Cont.)

Figures

Tables

Tables (Cont.)

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 in support of open science at Argonne National Laboratory and at Oak Ridge National Laboratory and deploys diverse petascale machines that are among the most powerful systems in the world today. Strategically, the LCF ranks among the top U.S. 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 around the clock every day of the 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 U.S. 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.

As of 2017, the ALCF now operates two production petascale resources for the research community: an Intel-based Cray system, Theta, and its IBM system, Mira. Since going online last July, Theta has taken on a diverse workload that includes *ASCR Leadership Computing Challenge (ALCC)* projects, *Director's Discretionary (DD)* projects, and *ALCF Data Science Program (ADSP)* projects. Meanwhile, Mira delivered 4.5 billion core-hours to 30 *Innovative and Novel Computational Impact on Theory and Experiment (INCITE)* projects and 1.6 billion core-hours to ALCC projects (29 awards for the 2016–2017 ALCC year and 22 awards for the 2017–2018 ALCC year); and a wide range of DD projects. Yet again, as Table ES.1 shows, Mira had an excellent year in terms of overall availability (96.4 percent), scheduled availability (99.5 percent), and utilization (95.5 percent; Table 2.1). Theta's performance statistics were strong, too, in terms of overall availability (90.0 percent), scheduled availability (98.6 percent), and utilization (82.7 percent; Table 2.1).

Moreover, ALCF's user community published more than 200 papers in high-quality, peer-reviewed journals and technical proceedings. The ALCF also offered a comprehensive program of high-performance computing (HPC) support services to help its community make productive use of the facility's diverse and growing collection of resources.

As the LCF prepares to enter the exascale era in the coming years, ALCF researchers are already leading and contributing to several strategic activities that aim to push the boundaries of what's possible in computational science and engineering. When the 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 Current Year (2017) Metrics

Area	Metric	2017 Target	2017 Actual
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	4.6/5.0
	User Survey – **Response Rate**	25.0%	46.9%
	% User Problems Addressed Within Three Working Days	80.0%	95.8%
Business Results	Theta Overall Availability	80.0%	90.0%
	Theta Scheduled Availability	90.0%	98.6%
	Mira Overall Availability	90.0%	96.4%
	Mira Scheduled Availability	90.0%	99.5%
	% of INCITE core-hours from jobs run on 16.7% or more of Mira (131,072–786,432 cores)	40.0%	72.2%
	% of INCITE core-hours from jobs run on 33.3% or more of Mira (262,144–786,432 cores)	10.0%	27.5%

Section 1. User Support Results

Are the processes for supporting the users, resolving users' problems, and conducting outreach to the user population 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 2017 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) 2017.

Table 1.1 All 2017 User Support Metrics and Results [a]

		2016 Actual	2017 Target	2017 Actual
Number Surveyed		1,019	N/A	976
Number of Respondents (Response Rate)		460 (45.1%)	25.0%	458 (46.9%)
Overall Satisfaction	Mean	4.5	3.5	4.5
	Variance	0.6	N/A	0.5
	Standard Deviation	0.7	N/A	0.7
Problem Resolution [b]	Mean	4.5	3.5	4.6
	Variance	0.4	N/A	0.5
	Standard Deviation	0.6	N/A	0.7
User Support [b]	Mean	4.5	3.5	4.5
	Variance	0.3	N/A	0.5
	Standard Deviation	0.6	N/A	0.7
		2016 Actual	**2017 Target**	**2017 Actual**
% User Problems Addressed Within Three Working Days [c]		95.9%	80.0%	95.8%

[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 2017 user survey was managed by Marketing Synergy, Inc., and developed in collaboration with ALCF staff. An effort was made to shorten the 2017 survey while retaining questions required for this 2017 Operational Assessment Report (OAR). The team also incorporated lessons learned from previous surveys and feedback from ALCF leadership and various ALCF teams. The final survey consisted of 24 questions, down from 43 questions in the previous year, and took approximately 10 minutes to complete. The survey e-mail campaign commenced on November 19, 2017, and consisted of an initial invitation to participate, followed by customized reminders sent out every Sunday and Thursday until the minimum response rate was reached in mid-December. Each reminder e-mail contained a message authored by various members of ALCF's leadership and a link to the online survey. Once a participant completed the survey, his or her e-mail address was removed from the distribution list for all subsequent reminders.

Likert Scale and Numeric Mapping

Almost all Likert Scale questions in the ALCF user survey use a six-choice scale for rating user responses. This is a standard for surveys because (1) it is easy for users to quickly place the response to a question within a range of options; (2) it can be mapped to a numeric scale; and (3) given a certain sample size, it can be used with a normal distribution to obtain useful statistical results. The method allows for use of off-the-shelf statistics functions to determine variance and standard deviation.

ALCF follows a standard practice and maps the Likert Scale in this way or similar:

Statement	Numeric
Strongly Agree	5
Agree	4
Neutral	3
Disagree	2
Strongly Disagree	1
N/A	(No Value)

Some questions were not well suited to the six-point scale and a different scale was therefore used. The only question included as part of the OAR to which this different treatment applied was the overall satisfaction question. It used the following five-choice scale:

Statement	Numeric
Excellent	5
Above Average	4
Average	3
Below Average	2
Poor	1

Some non-metric questions were revised on the 2017 User Survey to capture sentiments about various aspects of ALCF's user services that used the options below:

Select all that apply.
Praise
Suggestions for Improvement
Average
Below Average
Poor

Comments

1.1 User Support Metrics

A total of 976 individuals met the definition of a facility user and were invited to complete a user survey. Of these, 458 responded, for a 46.9 percent response rate — far superior to a generally accepted standard for survey response rates of 10 percent for this size of population. ALCF surpassed all targets for the survey metrics.

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

Table 1.2 2017 User Survey Results by Allocation Program

2017 Metrics by Program		INCITE	ALCC	INCITE + ALCC	DD	All
Number Surveyed		366	215	581	395	976
Number of Respondents		184	105	289	169	458
Response Rate		50.3%	48.8%	49.7%	42.8%	46.9%
Overall Satisfaction	Mean	4.7	4.5	4.6	4.5	4.5
	Variance	0.3	0.6	0.4	0.6	0.5
	Standard Deviation	0.5	0.8	0.6	0.8	0.7
Problem Resolution	Mean	4.7	4.6	4.6	4.7	4.6
	Variance	0.5	0.6	0.5	0.7	0.5
	Standard Deviation	0.7	0.7	0.7	0.8	0.7
User Support	Mean	4.6	4.5	4.6	4.6	4.5
	Variance	0.5	0.6	0.5	0.7	0.5
	Standard Deviation	0.7	0.8	0.7	0.9	0.7
All Questions	Mean	4.6	4.5	4.6	4.6	4.5
	Variance	0.5	0.8	0.5	0.7	0.5
	Standard Deviation	0.7	0.9	0.7	0.8	0.7

As Table 1.3 shows, in 2017, ALCF again exceeded the targets for overall satisfaction and user support.

Table 1.3 2016 and 2017 User Support Metrics

Survey Area	2016 Target	2016 Actual	2017 Target	2017 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

Table 1.4 shows the target set for the percentage of problem tickets addressed in three days or less, which the ALCF exceeded. A ticket is defined as “addressed” once the following conditions have been met: (1) the ticket is accepted by a staff member; (2) the problem is identified; (3) the user is notified; and (4) the problem is solved, or it is in the process of being solved.

Table 1.4 Tickets Addressed Metric

	2016 Target	2016 Actual	2017 Target	2017 Actual
% User Problems Addressed Within Three Working Days [a]	80.0%	95.9%	80.0%	95.8%
Average of Problem Resolution Ratings	3.5/5.0	4.5/5.0	3.5/5.0	4.6/5.0

[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 6,214 support tickets in 2017. The distribution of tickets across categories did not change significantly: The highest-percentage decrease in number of tickets was in the ‘Access’ category, from 18 percent in 2016 to 15 percent in 2017. The ‘Access’ change is statistically insignificant. The ‘Accounts’ category continued to be the largest (see Table 1.5). All of the other categories came within one percentage point of last year’s numbers.

In the third quarter, the user support team replaced the overarching ‘System’ ticket category with seven, more specific categories, as follows: ‘Compilers,’ ‘Debugging and Debuggers,’ ‘Libraries,’ ‘Network,’ ‘Performance and Performance Tools,’ ‘Reports,’ and ‘Scheduling.’ Furthermore, the ‘I/O and Storage’ and ‘Quota Management’ categories were replaced with ‘HPSS and Quota Management’ and ‘File System,’ respectively. This change allows the ALCF to better understand the types of tickets it fields and where improvements are needed.

Comparing the previous ‘System’ category, which represented eight percent of total tickets in 2016 (or 467 of 6,018), with the new 2017 categories and the ‘System’ category for the first two quarters of 2017, that percentage is 9 percent. This result shows that the new categories are statistically equivalent to the 2016 categories.

Comparing the number of tickets in the 2016 categories for ‘I/O and Storage’ (184) and ‘Quota Management’ (27), the percentage was 3.5% of the total tickets. In the 2017 categories ‘HPSS and Quota Management’ (60) and ‘File System’ (81) and adding in the tickets for the first half of the year for the original categories (‘I/O and Storage’ [103] and ‘Quota Management’ [25]), the percentage of tickets was 4 percent. Therefore, the new categories are statistically equivalent to the old categories in terms of percentages.

Table 1.5 Ticket Categorization for 2016 and 2017

Category	2016	2017 (Q1, Q2)	2017 (Q3, Q4)	2017 Total
Access	1,058 (18%)	476 (16%)	447 (14%)	923 (15%)
Accounts	2,755 (46%)	1,382 (45%)	1,488 (47%)	2,870 (46%)
Allocations	603 (10%)	326 (11%)	268 (8%)	594 (10%)
Applications Software	179 (3%)	88 (3%)	88 (3%)	176 (3%)
Automated E-mail Responses	480 (8%)	219 (7%)	341 (11%)	560 (9%)
Compilers	N/A	N/A	52 (2%)	52 (1%)
Data Transfer	41 (1%)	11 (0%)	27 (1%)	38 (1%)
Debugging and Debuggers	N/A	N/A	12 (0%)	12 (0%)
File System	N/A	N/A	81 (3%)	81 (1%)
HPSS [a] and Quota Management	N/A	N/A	60 (2%)	60 (1%)
Libraries	N/A	N/A	18 (1%)	18 (0%)
I/O [a] and Storage	184 (3%)	103 (3%)	N/A	103 (2%)
Miscellaneous	212 (3%)	125 (4%)	72 (2%)	197 (3%)
Network	N/A	N/A	6 (0%)	6 (0%)
Performance and Performance Tools	N/A	N/A	12 (0%)	12 (0%)
Reports	N/A	N/A	70 (2%)	70 (1%)
Quota Management	27 (0%)	25 (1%)	N/A	25 (0%)
Scheduling	N/A	N/A	107 (3%)	107 (2%)
System	467 (8%)	295 (10%)	N/A	295 (5%) [b]
Visualization	12 (0%)	6 (0%)	9 (0%)	15 (0%)
TOTAL TICKETS	6,018 (100%)	3,056 (100%)	3,158 (100%)	6,214 (100%)

[a] HPSS = high-performance storage system; I/O = input/output.

[b] 'System' was broken up into seven new categories for 2017 3Q and 4Q. Adding the 1Q/2Q 2017 'System' tickets to the other seven categories, the percentage is 9 percent.

1.3.1.2 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 weekly to discuss improvement opportunities and strategies, and to review any implemented changes.

The process to request and obtain an ALCF account has been streamlined in the following ways:

Accounts Website Changes

The ALCF continually strives to improve the overall quality and security of its user data, and of the user experience on its accounts website. This year, the following improvements were made to the ALCF accounts website to improve data collection and security:

- To curtail the number of ineligible account requests that ALCF receives, a new 'create account' webpage was implemented that requires a valid (eligible) e-mail address in order to proceed with an account request. In addition, a 'save for later' feature was introduced, which saves user-entered data for future completion.
- To lock down our user data, all administrative screens into our system were moved behind the laboratory's firewall. The administrative functions are no longer accessible to the general Internet.
- A science field drop-down list was added to the project creation process. This change facilitates improvement in reporting on projects.
- The field labeled "permanent address" on account request and account update screens was relabeled as "shipping address." This change facilitates better data entry into our user management system and clarity about the field itself.

Tracking Master User Agreements and User Acknowledgements

The team implemented a new, web-based solution for tracking Master User Agreements (MUAs) and user acknowledgements. These data were migrated from a legacy spreadsheet into a MySQL database, and a web interface was developed with fields for entering and retrieving information. The team tested the application and modified its workflows to adapt to the new interface. This application has improved process efficiency, as each team member can now access and edit real-time data. This improvement has directly impacted the account processing time for end-users.

Collaboration to Improve Collection of Project Quarterly and Annual Reports

The workflow for collecting ALCC quarterly and annual reports was revised and partially automated this year to reduce the number of staff members involved. This change was transparent to users.

1.3.2 Application Support

Modifying the LAMMPS code for Aurora

LAMMPS (Large-scale Atomic/Molecular Massively Parallel Simulator) is a classical molecular dynamics code and the workhorse of many projects at the Oak Ridge Leadership Computing Facility (OLCF), the ALCF, and other DOE leadership-class systems. In addition, LAMMPS serves as a simulation benchmark code for future CORAL systems, including CORAL-2 systems, and provides performance projections for Argonne's future exascale-class system, Aurora, expected in 2021. ALCF staff member Chris Knight made improvements to LAMMPS that drastically reduce the setup time. Previously, prior to a simulation run, information about specific atom bonding had to be scanned to compute the inter-atomic forces efficiently. LAMMPS's outdated communication patterns, which demand global inner-loop synchronizations, have become inefficient with large numbers of particles and processes. These were replaced with smarter, localized communication patterns. As a result, the setup time on the whole of Mira for a protein system containing 36.86 billion particles was reduced from a projected time of 18 hours to just 42 seconds. Similarly, the setup time on Mira for the calculation of electrostatic interactions on a system with 37 billion particles was reduced from an estimated 12 days to less than 1 second.

Resolving Memory Limitations for I/O Buffers in c64 Mode

In support of an INCITE allocation (Principal Investigator [PI]: Lars Bildsten), ALCF staff member Graham Fletcher worked with project team member Yan-Fei Jiang to solve a problem of insufficient memory for I/O buffers in c64 mode. Graham initially suggested three options: (1) fall back to 32 MPI (Message Passing Interface) ranks per node (c32 mode); (2) use multi-threading (OpenMP) to have a single MPI rank per node; or (3) pass all output data through a single rank on each node given that there is a single I/O channel per node. Yan-Fei rejected option 1 because the team would forfeit significant performance gains. Yan-Fei then considered option 3, because, while multi-threading a scientific code can be a significant task, his team had considered it a medium-term goal for the application. He ultimately selected the second option because the issue provided motivation toward the goal of multi-threading Athena++. The memory problem is now solved.

MPI-Driver Design Manages Calculations for Materials Science

In support of an ALCC allocation (PI: J. Ilja Siepmann), ALCF staff member Chris Knight worked closely with the project team to port its code, MCCCS-MN, to Mira. Materials discovery can involve sampling hundreds of thousands of candidates — a level of throughput that can only be achieved with the efficient exploitation of supercomputer resources. Chris helped design an MPI-based driver to manage large ensemble calculations for more than 330,000 zeolite structures. Work involved addressing IBM XL compiler issues, performance profiling, the addition of OpenMP parallelism to reduce memory requirements, optimizing I/O by reducing the number of output files, and developing workarounds to overcome limitations in both the Blue Gene/Q operating system and the job scheduler. As a result, MCCCS-MN has the flexibility to run on all of Mira, or in a "best-fit" mode to maximize throughput on Mira's queues.

HOPping from Cray/SGI to IBM Blue Gene/Q with High Efficiency

In support of an INCITE allocation (PI: Jonathan Poggie), ALCF staff member Ramesh Balakrishnan helped port the HOPS code to Mira from a Cray/SGI platform, work that included converting the build to use the IBM XLF compiler instead of an Intel compiler. The large I/O requirements of HOPS also had to be addressed, requiring a careful consideration of all of the I/O functionality available on Mira. For example, Ramesh used MPI codelets to show how certain features of Fortran 2008-derived data types, which are not directly supported by the XLF compiler, could be implemented to allow more efficient data transfer on the Blue Gene/Q. In addition, the MPI_Info interface can be used to inform the MPI library regarding the use of collective read/write operations, depending on the file sizes and typical usage patterns of computational fluid dynamics (CFD) codes. Under Ramesh's guidance, HOPS now efficiently scales to the full size of Mira.

Optimizing VSVB Kernel, Reducing Time to Solution Submitted for the Association for Computing Machinery (ACM) Gordon Bell Prize

In support of an INCITE allocation (PI: Maria Chan), ALCF staff member Graham Fletcher identified the kernel of VSVB (Variational Subspace Valence Bond) as a potential candidate for optimization. From February through April, Graham worked with Margaret Butler Fellow

Colleen Bertoni, on-site Intel representative Michael D'Mello, and ALCF performance engineer Vitali Morozov to improve the performance of VSVB on Mira and Theta. The VSVB kernel involves a single loop that evaluates Givens rotations. This kernel was optimized and substantially revised to incorporate a new method for handling up to three rotations in each loop iteration, whereas previous methods only evaluated a single rotation per loop iteration. By reusing data common to the rotations, the data traffic to and from the L1 cache memory has been significantly reduced. Overall, the time to solution was reduced by 5×, achieving one quarter of peak floating-point operations (FLOPs) on both platforms. These impressive gains were written up as a submission to the ACM Gordon Bell Prize for 2017.

A Director's Discretionary Team Reaches Its Goal of Obtaining a 2018 INCITE Award

In support of a DD allocation (PI: Allen Taflove), ALCF staff member Wei Jiang worked with a project team to port its application to Mira and to submit a successful 2018 INCITE proposal. As part of the porting process, Wei traced the failure of FDTD (finite-difference time-domain) during initialization on Mira to the use of the MPI_IN_PLACE argument together with MPI collectives on the Blue Gene /Q architecture and implemented a temporary send/recv buffer in FDTD that fixed the problem. He also helped to debug an input format problem on Mira by converting the binary storage format from little endian to big endian. Furthermore, Wei advised the code developer on how to implement a scalable initialization with MPI I/O functions and the PI's team on how to prepare capability benchmarks and run applications with optimal runtime environment variables. With Wei's help, the team was able to produce the PI's first INCITE proposal.

Reducing Core Hours at a Significant Pace – Accelerated Climate Modeling for Energy

In support of an INCITE allocation (PI: Mark Taylor), ALCF staff member Elise Jennings worked with the ACME (Accelerated Climate Modeling for Energy), now called the Energy Exascale Earth System Model (E3SM), team. The team's project goal is to achieve a high-resolution, fully coupled 50-year simulation. The components of a typical climate simulation — atmosphere, land, ocean, and sea ice — have greatly varying loads demanding different levels of parallelism. At the same time, some components can be run concurrently, whereas others must be run sequentially. Organizing all of these components to run efficiently on a supercomputer requires fine-tuning a static load balancing scheme across all of the model components. Elise worked with a project team member and Argonne's Mathematics and Computer Science (MCS) division postdoc Azamat Mametjanov to benchmark and improve code performance. Their work targeted the OpenMP threading and scheduling in the sea ice module together with the 'coupler' module that schedules tasks between the different model components. The per-thread load within each of the many sections of the two modules was optimized. As a result of their efforts, the cost has been reduced from 8 million core-hours to 7.1 million core-hours per simulated year, when running on 8,192 nodes with 64 OpenMP threads per node for a total of more than half a million threads. Work is ongoing.

Hydrodynamics into HACC to Enable Full Theta Simulation

In support of an Early Science Program (ESP) allocation (PI: Katrin Heitmann), ALCF staff member Adrian Pope worked with Theta ESP postdoc Jeffrey Emberson to incorporate a hydrodynamics capability into the Hardware Accelerated Cosmology Code (HACC) for execution on Intel many-core architectures. Adrian and Jeffrey improved the code in two important ways. First, they increased the efficiency of particle caches, necessary to handle the increased data traffic in the hydrodynamics functionality. Second, they addressed a build-up of floating point error that was found to occur in that code. Adrian also worked with ALCF performance engineer Vitali Morozov and the Intel compiler team to optimize the short-range gravity code in HACC to enable hydrodynamics simulations on the whole of Theta.

1.3.3 Resource Support

1.3.3.1 General Support

Jupyter Notebook

Jupyter Notebook is an interactive web application that allows users to create notebooks (or documents) containing active code, text, and visualization. This tool provides an intuitive interface for learning to code and to perform data analysis. The ALCF supports a proof-of-concept Jupyter service, accessible through crypto, and provides users with a Jupyter notebook and access to Mira's filesystems (/home, /projects) for data analysis. The service also offers limited job submissions to visualization cluster Cooley through Jupyter's terminal interface.

Condor

ALCF staff deployed a pilot of HTCondor CE (compute element) for Mira and Cooley, with custom integration scripts for the Cobalt job scheduler. The HTCondor CEs allow authenticated users to submit jobs remotely to ALCF systems without establishing an interactive shell login. HTCondor is a job workflow management system widely used in the high-energy physics community and by organizations such as the Open Science Grid. An HTCondor interface provides a standardized interface to ALCF systems and opens up the possibility of easy integration with other external workflow management systems in the future. Other possible future uses include data transfer workflows and co-scheduling between ALCF systems. The ALCF plans to deploy a similar HTCondor integration for Theta next year. The ATLAS project has performed initial pilot testing on Cooley with the goal of developing an integration with its own internal workflow management system using standardized components that can be easily deployed at other computing sites. This effort is expected to gain traction after the Theta deployment is complete. A second pilot project at Fermilab (HEPCloud) is starting initial testing for workflows targeting Mira.

1.3.4 Outreach Efforts

1.3.4.1 General Outreach

Facility tour requests are important outreach opportunities to inform the public about ALCF activities. The ALCF's process for managing tour requests begins with a form, which gathers basic information about the individuals or group requesting the tour, including group type, tour type, etc. The ALCF tour manager categorizes the request as general or VIP and schedules a suitable staff member to lead the tour. Rather than following a rote script, each guide speaks from his or her individual expertise and experience. In 2017, ALCF provided tours for approximately 45 groups and more than 400 people.

User Advisory Council

The ALCF's User Advisory Council (UAC) meets throughout the year and provides valuable input on key technology upgrades, messaging and communication, and user-centric metrics. The UAC is composed of 6–7 PIs with active projects from across the three allocation programs (INCITE, ALCC, DD). In 2017, UAC members evaluated and provided feedback on the following ALCF activities: training activities related to Intel's Knights Landing (KNL) architecture; the 2016 ALCF user survey; booting into different memory modes on Theta; Theta's release into production; scheduling policies; and matters related to Aurora. The UAC also served as a sounding board for communication strategies related to the allocation programs.

Connection to Technology Commercialization and Partnerships Division

The ALCF has forged a strong relationship with the Technology Commercialization and Partnerships (TCP) division in Argonne's new Science and Technology Partnerships and Outreach directorate. The ALCF frequently supports meetings that TCP organizes with potential industry partners. While many of these meetings do not focus on high-performance computing (HPC), involving the ALCF can lead to collaborations later on. Conversely, ALCF involves TCP in HPC-specific conversations with industry so that the potential partner can obtain a more complete picture of Argonne capabilities and partnership opportunities.

Support of GlobalFoundries ALCC Extension

Sefa Dag from semiconductor company GlobalFoundries was awarded an ALCC allocation in June 2016; however, because of concerns about Argonne's MUA, Dag was unable to start work. The ALCF worked with GlobalFoundries and Argonne's legal department to address the company's concerns, and an agreement was signed in May 2017. ALCF communicated the issues to the DOE ALCC program manager and worked out a plan to provide Dag with an extension through December 2017.

1.3.4.2 Workshops, Webinars, and Training Programs

The ALCF conducts a range of workshops, webinars, and training programs to support the efforts of its users. 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

The ALCF hosts several virtual 'Getting Started' sessions throughout the year, oriented to teams and small groups and conducted via videoconference. These sessions are taught by ALCF staff and cover system overviews, code building, storage, operating system (OS) and filesystems, compilers, tools, queues, hands-on exercises, and other topics.

Preparing for KNL — Videos for Users

Intel and the ALCF curated a collection of publicly available videos aimed at getting users up and running on KNL hardware. The ALCF conducted an e-mail campaign to promote each new video and tracked the e-mail statistics (views, clicks, etc.).

Preparing for KNL — Developer (Jam) Sessions

The ALCF organized a three-part, videoconference seminar series that convened experts from Intel and Cray and early Theta users — project team members from the Theta ESP and ALCF Data Science Program (ADSP) — to discuss and address specific KNL-related issues encountered on Theta.

ALCF Computational Performance Workshop

The ALCF held a four-day scaling workshop in May. Formerly known as the "Scaling Your Science on Mira" workshop, this annual event is a 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 majority of this year's workshop was devoted to hands-on application tuning on Mira and Theta, and talks covered a range of topics, including ensemble jobs, parallel I/O, data analysis, and data science. More than half of the 60+ attendees stated a pre-event goal of submitting a new or renewal 2018 INCITE project proposal. Roughly 40 percent of the attendees who submitted a proposal were later awarded an INCITE or ALCC allocation.

Theta — Many-Core Tools and Techniques and Developer Sessions

To help support and expand the Theta user base, ALCF introduced a tutorial series and a live-presentation webinar series. Topics covered so far in the "Many-Core Tools and Techniques" tutorial series include architectural features, tuning techniques, analysis tools, and software libraries. The "Many-Core Developer Sessions" webinar series was created to foster discussion between actual developers of the emerging many-core hardware and software, and the early users of that technology. Speakers in this series have included developers from Intel and Allinea (ARM), covering topics such as AVX-512, MKL, Vtune, TensorFlow, and debugging on KNL. Future speakers include developers from Cray and ParaTools. The format of these webinars includes a 30-minute presentation, after which audience participation is encouraged. The ALCF has hosted seven webinars to date.

ATPESC 2017

Last July and August, ALCF ran its fifth Argonne Training Program on Extreme-Scale Computing (ATPESC), a highly competitive, two-week training program for 65–70 early career researchers and aspiring computational scientists. ATPESC's seven program tracks focus on HPC methodologies that are applicable to current and future machine architectures — including exascale systems. Renowned scientists, HPC experts, and leaders provide the technical lectures, hands-on exercises, and dinner talks. To further extend ATPESC's reach, ALCF staff produces and uploads video playlists of ATPESC content to Argonne's YouTube training channel. More than 60 hours from the 2017 course were uploaded in September, and promotion is ongoing. Since 2013, ATPESC's inaugural year, these videos have been viewed more than 30,000 times. In 2017, ATPESC organizers introduced several improvements to the logistics and management of the event, including using the Box file sharing service to collaborate with the audio/video company contracted to record the sessions.

1.3.4.3 Community Outreach

Youth Outreach

Women in STEM

Every February, Argonne hosts Introduce a Girl to Engineering Day (IGED) for local eighth grade girls and pairs them with Argonne mentors to participate in group activities focused on STEM (science, technology, engineering, and mathematics) careers. ALCF staff member Jini Ramprakash serves on the organizing committee, and many other ALCF volunteers serve as mentors or as activities supervisors. ALCF staff member Liza Booker was named co-chair of IGED 2018 and commenced work for the 2018 event in September 2017.

ALCF staff members contribute to many other STEM-focused events aimed at women and girls during the year, including: Women in Science and Technology (WIST), Argonne's Science Careers in Search of Women conference (SCSW), AnitaB.org (formerly the Anita Borg Institute); and the Grace Hopper Celebration of Women in Computing.

Hour of Code

As part of Code.org's annual event, several ALCF staff members visited various Chicago and suburban classrooms during Computer Science Education Week (CSEdWeek), which took place December 4–10, 2017, where they gave talks and demos and led Hour of Code activities. CSEdWeek was established by Congress in 2009 to raise awareness about the need to elevate computer science education at all levels.

Summer Coding Camp

ALCF staff taught and mentored 30 area high school students at Argonne's Summer Coding Camp from July 24–28, 2017. The camp curriculum promotes problem-solving and teamwork skills through hands-on coding activities, such as coding with Python and programming a robot via a Raspberry Pi. The camp is a joint initiative of ALCF and Argonne's Educational Programs Office.

CodeGirls@Argonne Camp

Argonne hosted nearly two dozen seventh and eighth grade girls for a two-day camp held August 3–4, 2017, and aimed at teaching the fundamentals of coding in the Python programming language. The girls also met with female computer scientists at the lab and toured the ALCF.

Visits and Tours

As noted in Section 1.3.4.1, 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 give ALCF staff the opportunity to interact with a wide range of visitors having differing levels of expertise and varying interests. Below is a selection of the groups that toured the facility in 2017.

International Delegations

In January, ALCF hosted the chief financial officer of Sabanci Holding and the president of EnerjiSa, two Turkish entities that are active in a wide spectrum of energy technologies. In May, interns of the Trade Commission of Spain toured the ALCF. In July, UK Special Representative on Climate Change Nick Bridge visited the ALCF, accompanied by Afriqnmun Lovejoy, Head of International Climate Outreach, UK Department for Business, Energy & Industrial Strategy; Benjamin Springer, Senior Policy Advisor Energy and Environment, British Embassy Washington; and two members of the British Consulate-General Chicago. In November, a high-level nuclear engineering delegation from Indonesia, hosted by representatives from Thorcon Power, came through.

School Groups

In April, the ALCF hosted two STEM-oriented university student organizations: The University of Illinois at Urbana-Champaign Chapter of the Society for Industrial and Applied Mathematics (SIAM) and a group from Chicago State University. In May, ALCF hosted science and engineering graduate students from the University of Iowa. In August, undergraduates from Missouri Science & Technology's materials science and engineering department visited, as well as a group of 10th and 11th grade students from Huantai 1st Middle School in Zibo, China, and four high school student members of the Sirius Astronomy Association in Algeria. In November, ALCF welcomed the Computer Science Honors Society from North Central College in Naperville, Illinois.

Industry Groups

In June, the Illinois Chamber of Commerce Executive Board visited ALCF. In July, a review team from the DOE Office of Nuclear Technology Research and Development included an ALCF tour in their visit to Argonne.

Other STEM Activities

ALCF staff member Carissa Holohan attended the 7th Annual STEM Conference in Chicago, held on November 16–19, 2017. Holohan staffed the Argonne career fair table and attended workshops and panels on issues impacting the LGBTQA (lesbian, gay, bisexual, transgender, questioning, and asexual) communities in STEM fields.

ALCF staff member Ben Lenard gave an invited talk at Moraine Valley Community College's STEM Center, where he presented an overview of the ALCF computing environment and his responsibilities as a database administrator.

ALCF staff member Richard Coffey organized a STEAM (STEM + Arts) workshop as part of the Broader Engagement (BE) program at the annual SIAM Computational Science and Engineering Conference (SIAM CSE2017), held on February 27–March 3, 2017. The SIAM CSE BE program was supported by Argonne and others and aimed to create a more inclusive environment for underrepresented groups, such as women, minorities, and people with disabilities.

1.3.5 Communications

Communications through Mailing Lists and Social Media

ALCF provided information to users through several electronic communication channels, including direct e-mails, custom 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 2017 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 in a text-based format per user and channel preference.	As needed

Table 1.7 2017 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 in ATPESC.

Promotional Activities and Media Hits

ALCF published 44 science stories in 2017 on the ALCF website and in Newsbytes newsletters. All of these science stories are also channeled to Argonne's media and public relations team. Thanks to the efforts of Argonne's media and public relations staff 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," that is, a person has made a decision to run or not run the story. In 2017, the facility posted 74 such media hits to the website. The media team uses the 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 2017, Meltwater captured 132 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.

Table 1.8 Publications Designed for Print in 2017

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

Conclusion

Our users are at the forefront of all of ALCF's interactions. As a user facility, the ALCF strives to improve our user experience processes and to help our customers make the most of their allocation time on our resources. In 2017, improvements included working with ALCC and INCITE project teams to transform and optimize their scientific codes, as well as helping DD projects reach their scientific goals and obtain INCITE awards. The ALCF worked in partnership with other national laboratories to present sessions on our work at the annual meetings of major national scientific societies. The ALCF implemented new tools to help with data analysis and conducted a wide range of outreach events to teach best practices and help our users explore new technologies. Once again, the ALCF met or exceeded all of our ticket and survey targets and continues to function at the level expected of a leadership facility.

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				Theta (Cray XC40): 3240-node, 21K-core, 56 TB MCDRAM, 679 TB DDR4 [h]	
	CY 2016		CY 2017		CY 2017	
	Target	Actual	Target	Actual	Target	Actual
Scheduled Availability	90.0%	98.9%	90.0%	99.5%	90.0%	98.6%
Overall Availability	90.0%	94.9%	90.0%	96.4%	80.0%	90.0%
System MTTI	N/A [c]	10.00 days	N/A	10.37 days	N/A	6.94 days
System MTTF	N/A	40.55 days	N/A	36.44 days	N/A	30.30 days
INCITE Usage	3.5B	3.9B [d]	3.5B	4.5B [f]	N/A	N/A
Total Usage	N/A	6.0B [e]	N/A	6.4B [g]	N/A	724.1M
System Utilization	N/A	90.8%	N/A	95.5%	N/A	82.7%
Mira INCITE Overall Capability [a]	40.0%	76.4% [c]	40.0%	72.2% [e]	N/A	N/A
Mira INCITE High Capability [b]	10.0%	40.9% [c]	10.0%	27.5% [f]	N/A	N/A

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

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

[c] N/A = Not applicable.

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

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

[f] Usage includes 3.0M core-hours from Cetus production jobs.

[g] Usage includes 13.4M core-hours from Cetus production jobs.

[h] See Appendix A for the definition of the Theta production node count calculation.

Note: DDR4 = double data rate fourth generation; MCDRAM = multi-channel dynamic random access memory; RAM = random access memory; TB = terabytes.

ALCF Resources

During CY 2017, ALCF operated one INCITE production resource, Mira, and two ALCC production resources, Mira and Theta. Mira is a 49,152-node, 786,432-core, 10-petaflops (PF) Blue Gene/Q with 768 Tebibytes (TiB) of RAM. Mira mounts three general parallel file systems (GPFSs) with approximately 26.5 petabytes (PB) of usable space. Mira has an associated visualization and analysis cluster called Cooley. As of December 17, 2017, Theta is a 4,392-node, 281,088-core, 11.69-PF Cray XC40 with 892 TiB of RAM; however, for most of the year, it was a 3,624-node, 231,936-core, 9.65-PF Cray XC40 with 736 TiB of RAM. (Appendix A contains node calculation information.) Theta mounts one GPFS and one Lustre file system with approximately 8.6 PB of usable space. Both Mira and Theta have access to the facility-wide HPSS tape archive. ALCF operated two other Blue Gene/Q systems, Cetus and Vesta, and one other Cray XC40, Iota.

Cetus is a 4,096-node, 65,536-core Blue Gene/Q with 64 TiB of RAM. Cetus shares file systems with Mira. Vesta is a 2,048-node, 32,768-core Blue Gene/Q with 32 TiB of RAM. Vesta is an independent test and development resource and shares no resources with Mira or Cetus. Iota is a 44-node, 2,816-core Cray XC40 with 8.9 TiB of RAM. Like Vesta, Iota is an independent test and development resource and shares no resources with Theta.

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 has allowed Mira to continue to operate as designed while enabling Cetus to support a new class of leadership applications.

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.

Theta

2.1.1 Scheduled and 2.1.2 Overall Availability

Theta entered full production on July 1, 2017. In consultation with the DOE Program Manager, ALCF has agreed to a target of 80 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 Theta's availability results.

Table 2.2 Availability Results for Theta

Theta (Cray XC40) 3,240-node, 21K-core, 56 TB MCDRAM, 679 TB DDR4		
	CY 2017	
	Target (%)	Actual (%)
Scheduled Availability	90.0	98.6
Overall Availability	80.0	90.0

The remainder of this section covers significant availability losses, and responses to them, for both scheduled and overall availability data. Details on the calculations can be found in Appendix A.

Explanation of Significant Availability Losses

This section briefly describes the causes of major losses of availability for the period July 1, 2017, through December 31, 2017, as annotated in Figure 2.1.

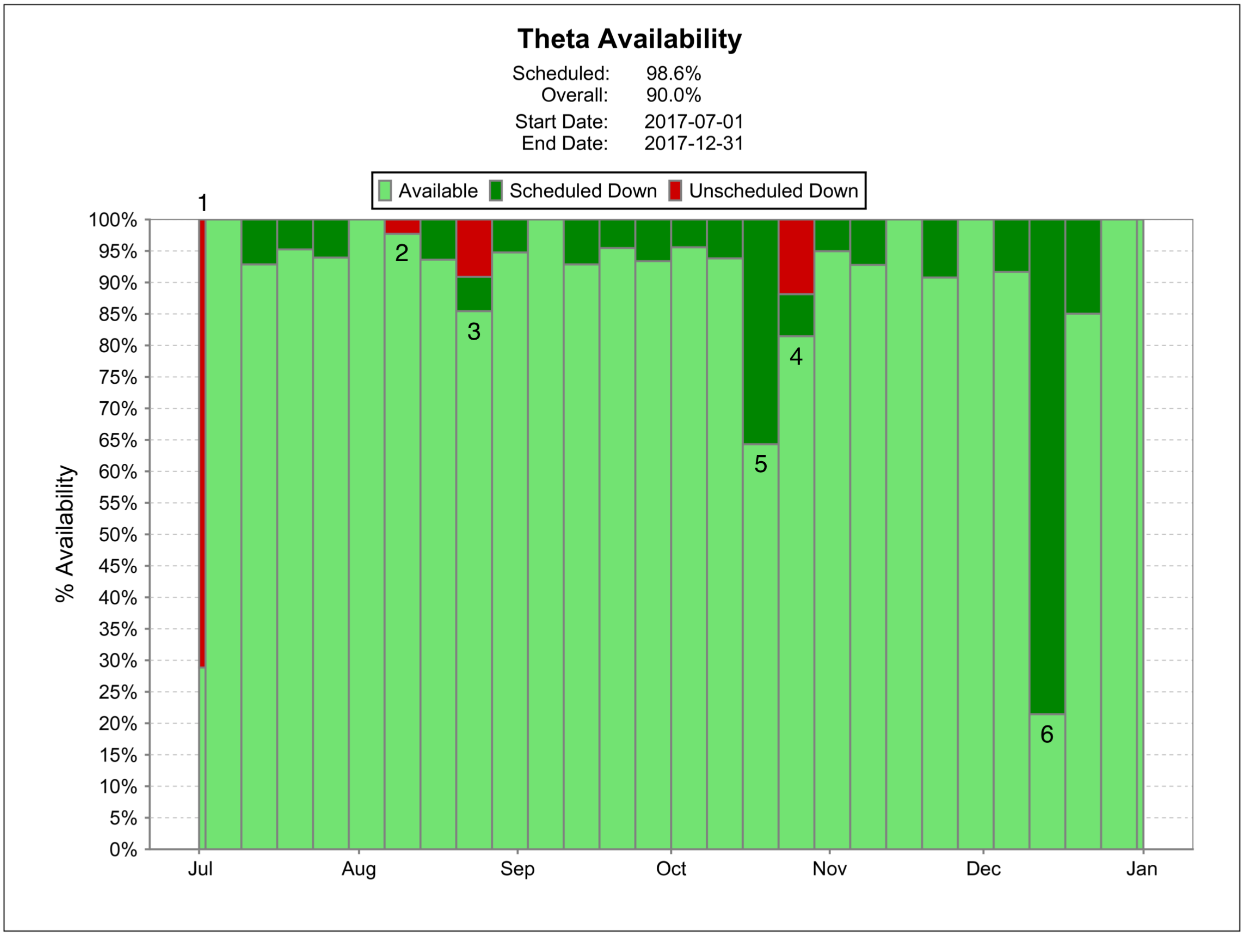

Figure 2.1 Theta's Weekly Availability for CY 2017

Graph Description: Except for the first and last bars, which indicate partial weeks, 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: Cobalt bug – July 1, 2017

Theta went in to production on July 1, 2017 at 00:00. July 1 was the last day of that week, so the reporting bar on the graph represents only one day. Shortly after going production, at 01:15, an interactive job triggered a Cobalt bug that prevented allocated nodes from being freed for use by subsequent jobs. The issue was corrected later that day around 18:15. A mitigation was deployed on July 3 to prevent the bug from being triggered until the final fix was applied on July 10 during scheduled maintenance.

Item 2: Lustre bug – August 15, 2017
A user went over quota, which triggered a Lustre bug (cl_sync_io_wait), and approximately 30 nodes went down as a result. A hotfix was applied the next maintenance day (August 21, 2017) while an official patchset was generated and later applied.

Item 3: Stale system state information – August 22, 2017
On August 22, the Application Level Placement Scheduler (ALPS), Cray's software suite for resource management, stopped releasing reservations because stale state information about one node was stored in the system state database. (ALCF's resource scheduler, Cobalt, interacts with ALPS to schedule jobs on Cray systems.) No new jobs started for about 8 hours overnight until this situation was discovered the next morning and state was manually cleared. It was ultimately determined that one down node was not marked as such in the state database. The node was manually marked down and the entire cluster rebooted to force a state database refresh. Once these steps were performed, jobs were able to run again. A bug report was filed with the vendor; however, this issue has not been encountered again.

Item 4: DDN failure – October 22, 2017, November 5, 2017
On October 22, a faulted disk in one of the storage enclosures (DDNs) triggered a known issue with the running Storage Fusion Architecture (SFA) OS version 2.3.1 affecting the mira-home filesystem. A reservation was put in place, and jobs that were already running were able to complete normally. The filesystem was restored as part of the normally scheduled maintenance the next day. This incident was the second occurrence of this bug, with the first occurring before Theta entered production; and a plan was developed to upgrade all of the DDNs to SFA/OS version 3.1.2 starting with our test and development Blue Gene/Q machine, Vesta, on November 6.

On November 5, this same SFA/OS bug was encountered again with the mira-home filesystem. Jobs again were held, and the affected DDN was rebooted, after which jobs were released to run again. On November 6, the Vesta DDNs were upgraded as previously planned, and testing was performed to help ensure stability and integrity. The mira-home filesystem was upgraded on December 18; however, because Theta did not undergo maintenance that day, its upgrade took place in January 2018 during the first joint maintenance between Mira and Theta.

Item 5: Software Update – October 16, 2017
Starting on October 16 and completing the following day, an operating system upgrade, UP04, was installed for the Cray Linux Environment (CLE) and System Management Workstation (SMW). Although the scheduled maintenance window was for three days, the ALCF was able to complete the upgrade in just two.

Item 6: Four Cabinet Expansion – October 16, 2017
Between December 11 and December 17, four additional XC40 cabinets (768 nodes) were delivered, installed, and accepted. A full acceptance test was performed including hardware and software diagnostics and verification, a successful performance run suitable for Top500 submission, and a suite of representative scientific applications.

2.1.3 System Mean Time to Interrupt (MTTI) and 2.1.4 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.

Theta MTTI and MTTF Summary

MTTI and MTTF are reportable values with no specific targets. Table 2.3 summarizes Theta's current MTTI and MTTF values.

Table 2.3 MTTI and MTTF Results for Theta

Theta (Cray XC40): 3,240-node, 21K-core, 56 TB MCDRAM, 679 TB DDR4		
	CY 2017	
	Target	**Actual**
System MTTI	N/A[a]	6.94 days
System MTTF	N/A	30.30 days

[a] N/A = Not applicable.

Theta currently functions on a weekly maintenance schedule. ALCF takes the machine out of service every Monday to perform Cray driver upgrades, hardware replacements, OS upgrades, etc. Further, while Theta is out of service concurrently with Mira, ALCF uses that opportunity to perform other potentially disruptive maintenance, such as facilities power and cooling work and storage systems upgrades and patching. Theta's weekly maintenance schedule caps MTTI at seven 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 Theta's utilization results, and Figure 2.2 shows Theta's utilization over time by program.

Table 2.4 System Utilization Results for Theta

Theta (Cray XC40): 3,240-node, 21K-core, 56 TB MCDRAM, 679 TB DDR4		
	CY 2017	
	Target	Actual
System Utilization	N/A [a]	82.7%

[a] N/A = Not applicable.

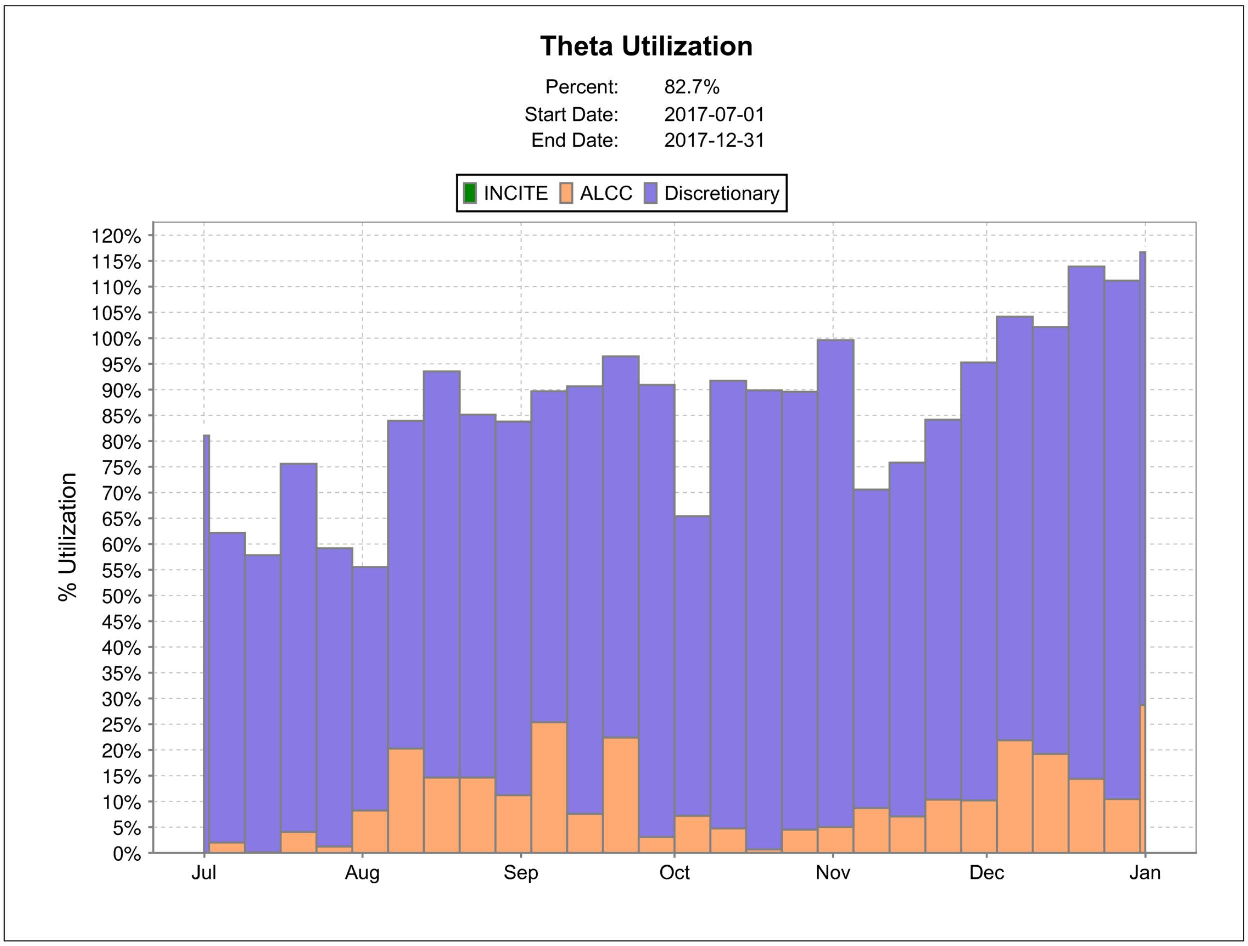

Figure 2.2 Theta's System Utilization over Time by Program

The system utilization for Theta was 82.7 percent for its 2017 production period of July 1, 2017, through December 31, 2017.

Table 2.5 shows how Theta'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. In 2017, Theta supported only two types of allocations: ALCC (30 percent) and DD (70 percent). 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. Normally, a majority of DD projects are exploratory investigations, so the time allocations are often not used in full; however, ESP allocations were also DD. DD allocations are discussed in detail in Section 3.3. In CY 2017, ALCF delivered a total of 724.1 million core-hours on Theta.

Table 2.5 Core-Hours Allocated and Used by Program on Theta

Theta (Cray XC40): 3,240-node, 21K-core, 56 TB MCDRAM, 679 TB DDR4			
	CY 2017		
	Allocated	Used	
	Core-hours	Core-hours	%
INCITE	N/A [b]	N/A	N/A
ALCC	209.5M	84.9M	11.7
DD [a]	628.2M	222.4M	30.7
ESP	603.0M	416.9M	57.6
Total	1.4B	724.1M	100.0

[a] Does not include ESP.

[b] N/A = Not applicable.

Summary: For CY 2017, the system usage and utilization values were in line with general expectations. The calculations for utilization are described in Appendix A.

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 Theta, capability is defined as using greater than 20 percent of the machine. See Appendix A for more detail on the capability calculation. Although Theta did not participate in the INCITE program for CY 2017, data are also provided in Table 2.6 for ALCC and DD projects as reference, and Figure 2.3 shows the overall distribution of job sizes over time on Theta while Figure 2.4 shows the capability hours used over time by program on Theta. No capability targets are set for ALCC and DD projects.

Table 2.6 Capability Results for Theta

Theta (Cray XC40): 3,240-node, 21K-core, 56 TB MCDRAM, 679 TB DDR4			
	CY 2017		
Capability Usage	**Total Hours**	**Capability Hours**	**Capability (%)**
INCITE Overall	N/A [b]	N/A	N/A
INCITE High	N/A	N/A	N/A
ALCC Overall	84.9M	37.6M	44.3
ALCC High	84.9M	1.0M	1.2
Director's Discretionary Overall [a]	222.4M	69.1M	31.1
Director's Discretionary High [a]	222.4M	26.4M	11.9
ESP Director's Discretionary Overall	416.9M	184.0M	44.1
ESP Director's Discretionary High	416.9M	107.7M	25.8
TOTAL Overall	724.1M	290.7M	40.1
TOTAL High	724.1M	135.1M	18.7

[a] Does not include ESP.

[b] N/A = Not applicable.

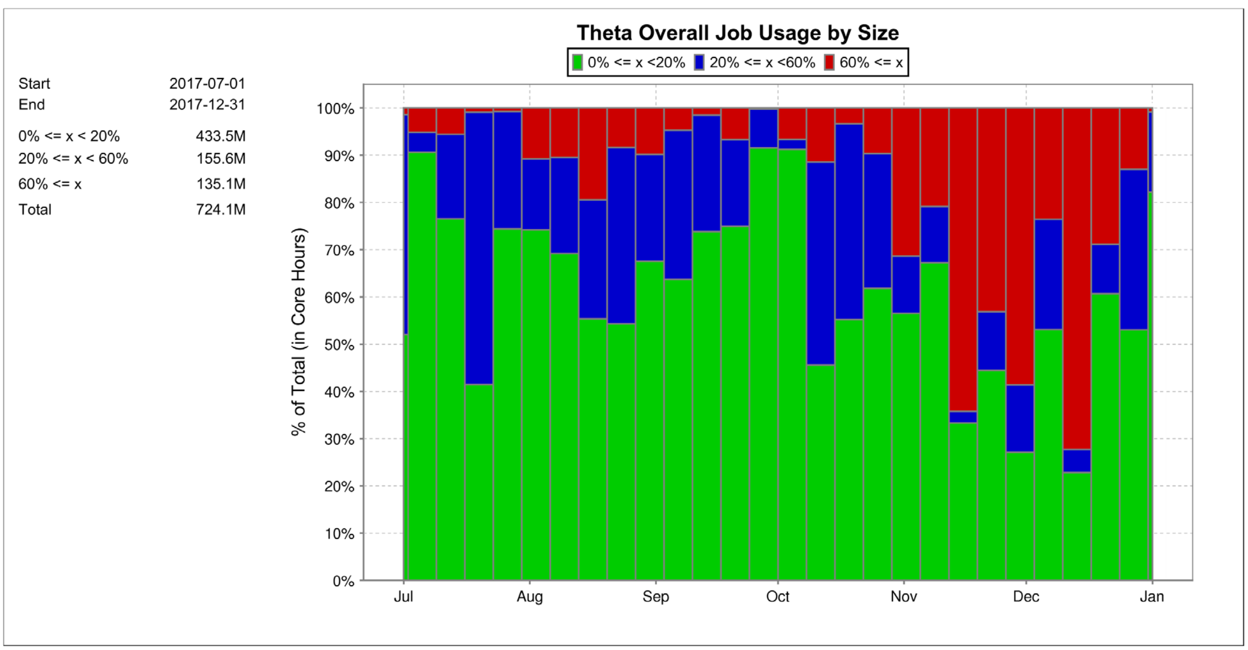

Figure 2.3 Theta's Job Usage by Size

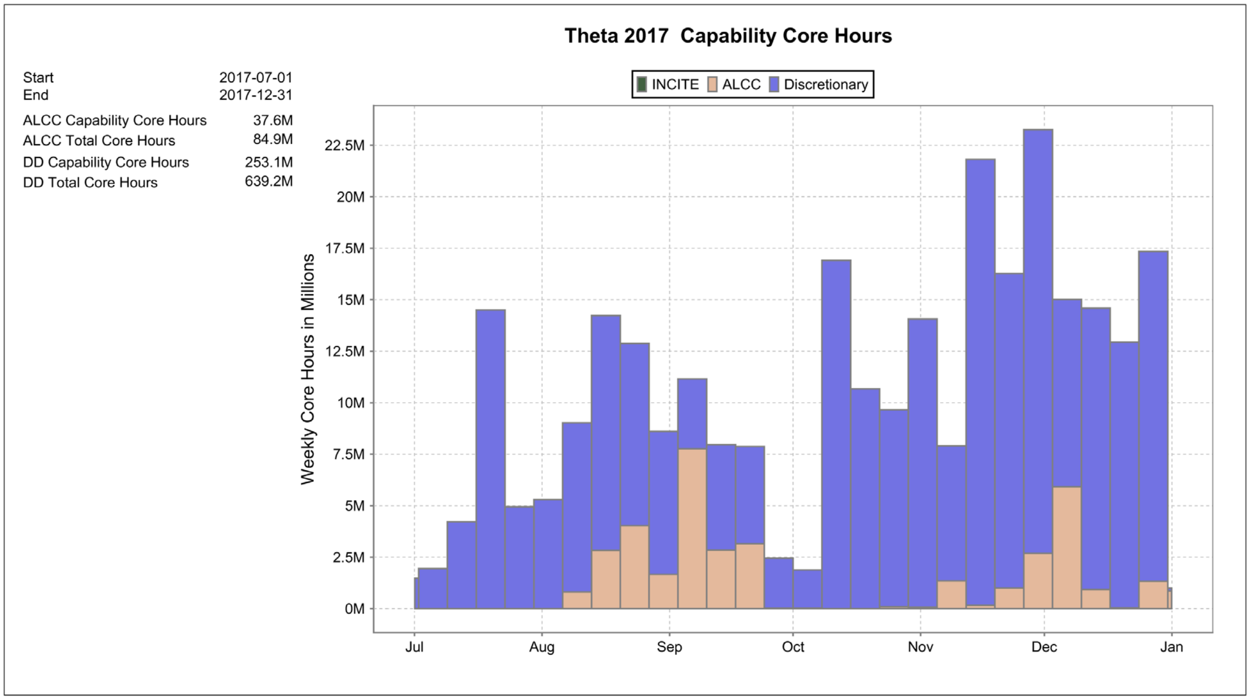

Figure 2.4 Theta's Capability Usage by Award Category

Mira

2.1.1 Scheduled and 2.1.2 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 2.7 summarizes Mira's availability results.

Table 2.7 Availability Results for Mira

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	CY 2016		CY 2017	
	Target (%)	Actual (%)	Target (%)	Actual (%)
Scheduled Availability	90.0	98.9	90.0	99.5
Overall Availability	90.0	94.9	90.0	96.4

The remainder of this section covers significant availability losses, and responses to them, for both scheduled and overall availability data. Details on the calculations can be found in Appendix A.

Explanation of Significant Availability Losses

This section briefly describes the causes of major losses of availability for the period January 1, 2017, through December 31, 2017, as annotated in Figure 2.5.

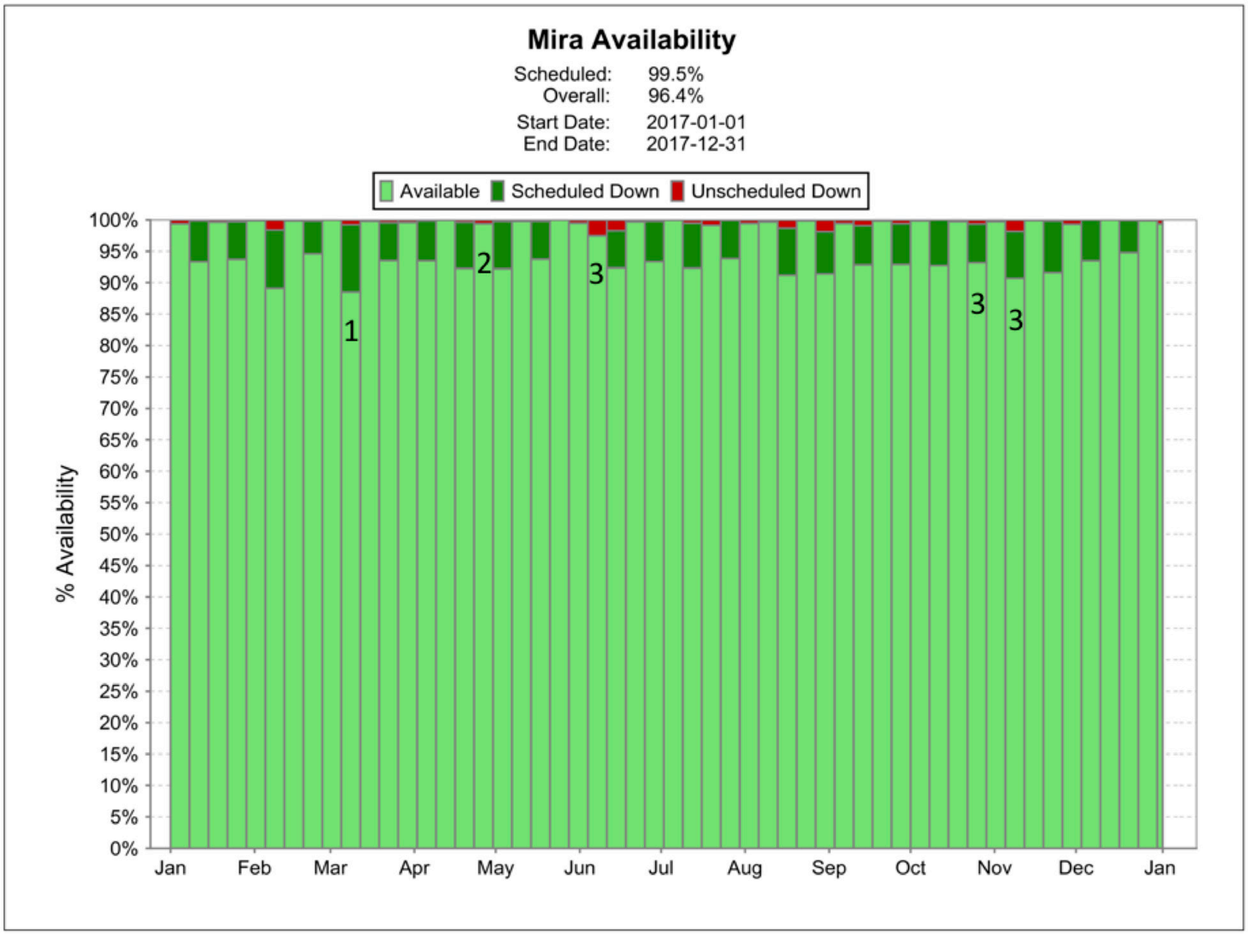

Figure 2.5 Mira's Weekly Availability for CY 2017

Graph Description: Except for the last bar, which accounts for only one day, each bar in Figure 2.5 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 5, 2017

On Sunday, March 5, mira-fs1 (a 7-PB scratch filesystem for Mira) hung. This failure 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. This bug was the same one encountered on March 13, 2016. We upgraded to GridScalar 3.2.1 over the months of April, May, and June during normal maintenance periods to fix the underlying bug.

Item 2: Subnet Manager failure – April 19, 2017
On Wednesday, April 19, the subnet manager on the Mira service node, mirasn2, received a `SIGBUS` signal from the *bgmaster* process, leaving behind evidence of an internal memory error, likely a memory leak. While the control system was able to recover from this initial error, high-availability (HA) monitoring lost its connection to the *bgagentd* process during subnet manager process recovery and subsequently began issuing restart requests. The *bgagentd* process acts as a message bus for a number of critical control system and job control processes and is vital to the operation of the control system. After a number of restarts, the *bgmaster* process forced a full-system restart of all components on the system resulting in the termination of all jobs including compute and ION partitions. After restarting and testing the HA monitoring and control system processes and rebooting the ION partitions, service was returned. Additional monitoring to identify a potential building memory problem with the control system components was implemented, as were changes to the HA configuration, ensuring that a failed monitoring thread would not force an agent restart.

Item 3: DDN failure – June 10, 2017, October 22, 2017, November 5, 2017
On Saturday, June 10, mira-fs1 operations started to hang. Once detected, all queued jobs using mira-fs1 were admin held and mira-fs0 jobs continued to run uninterrupted. The problem was determined to be a faulted disk in one of the DDNs that triggered a known issue with the running SFA/OS version 2.3.1. The affected DDN was rebooted, and the filesystem was returned and held jobs released.

Again, on October 22, a faulted disk in one of the DDNs triggered this SFA/OS bug—this time with the mira-home filesystem. A reservation was put in place, and jobs that were already running were able to complete normally. The filesystem was restored as part of the normally scheduled maintenance the next day. After this second incident, a plan was developed to upgrade all of the DDNs to SFA/OS version 3.1.2 starting with our test and development machine, Vesta, on November 6.

On November 5, this same SFA/OS bug was encountered again with the mira-home filesystem. Jobs again were held and the affected DDN was rebooted, after which jobs were released to run again. On November 6, the Vesta DDNs were upgraded as previously planned, and testing was performed to help ensure stability and integrity. The mira-fs1 filesystem was upgraded on November 20, and the mira-fs0 filesystem was upgraded on December 4. The mira-home filesystem was upgraded on December 18; as Theta did not undergo maintenance that day, it was upgraded in January 2018 during the first joint maintenance between Mira and Theta.

2.1.3 System Mean Time to Interrupt (MTTI) and 2.1.4 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.

Mira MTTI and MTTF Summary
MTTI and MTTF are reportable values with no specific targets. Table 2.8 summarizes the current MTTI and MTTF values for Mira.

Table 2.8 MTTI and MTTF Results for Mira

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	CY 2016		CY 2017	
	Target	Actual	Target	Actual
System MTTI	N/A [a]	10.00 days	N/A	10.37 days
System MTTF	N/A	40.55 days	N/A	36.44 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, OS upgrades, etc. Further, while Mira is out of service concurrently with Theta, 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 2.9 summarizes Mira's utilization results, and Figure 2.6 shows system utilization over time by program.

Table 2.9 System Utilization Results for Mira

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	CY 2016		CY 2017	
	Target	Actual	Target	Actual
System Utilization	N/A [a]	90.8%	N/A	95.5%

[a] N/A = Not applicable.

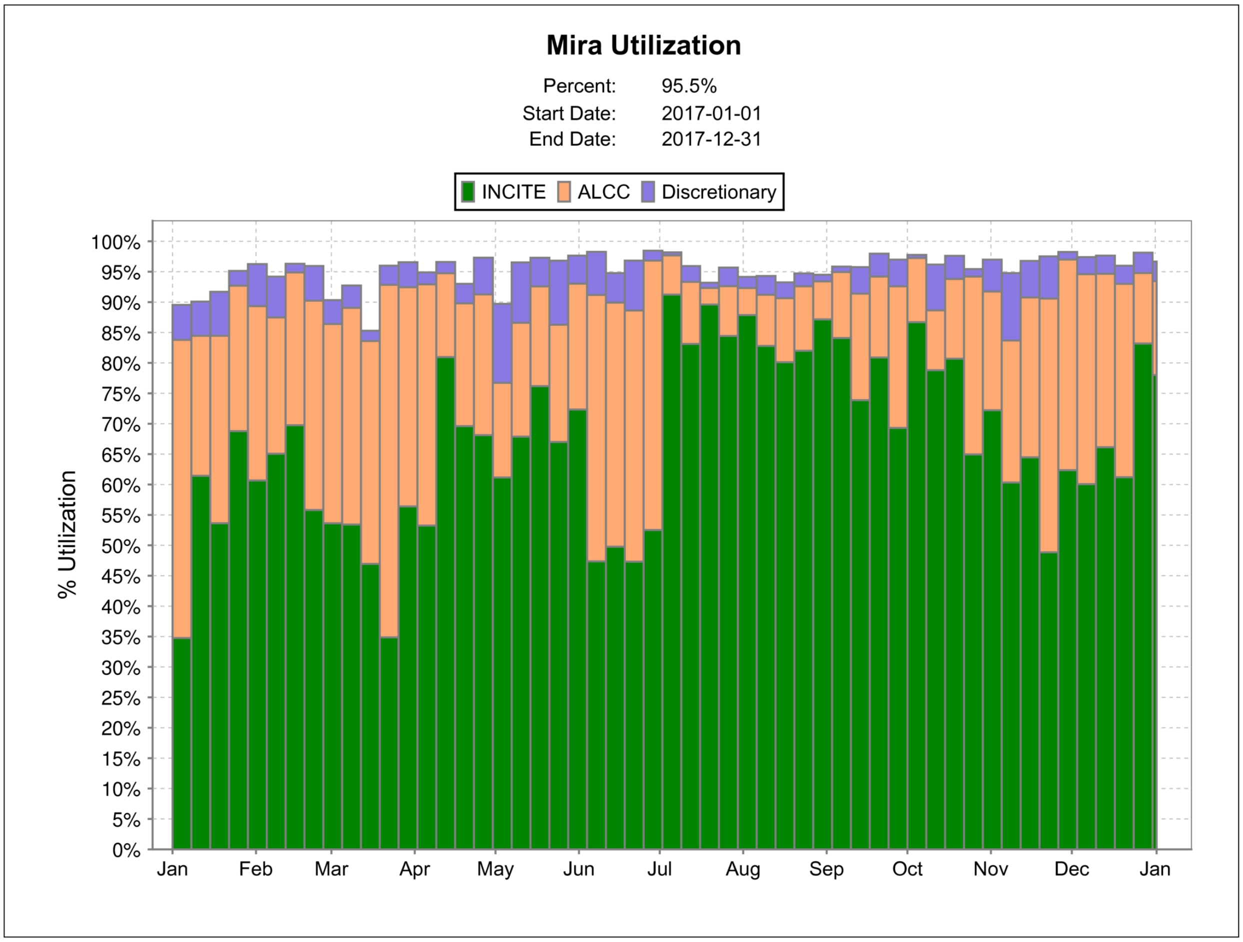

Figure 2.6 Mira's System Utilization over Time by Program

The system utilization for Mira was 95.5 percent for its 2017 production period of January 1, 2017, through December 31, 2017.

Table 2.10 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 are allocated to the INCITE program, up to 30 percent are available for the ALCC program allocations, and 10 percent are available for 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 2017, ALCF delivered a total of 6.4 billion core-hours on Mira.

Table 2.10 Core-Hours Allocated and Used by Program on Mira

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM						
	CY 2016			CY 2017		
	Allocated	Used		Allocated	Used	
	Core-hours	Core-hours	%	Core-hours	Core-hours	%
INCITE	3.5B	3.9B [a]	65.5	3.5B	4.5B [d]	70.3
ALCC	1.7B	1.7B [b]	28.4	1.8B	1.6B [e]	24.9
DD	1.0B	365.5M [c]	6.1	1.7B	309.6M [f]	4.8

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

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

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

[d] Usage includes 3.0M core-hours from Cetus production jobs.

[e] Usage includes 2.3M core-hours from Cetus production jobs.

[f] Usage includes 7.7M core-hours from Cetus production jobs.

Summary: For CY 2017, the system usage and utilization values were in line with general expectations. The calculations for utilization are described in Appendix A.

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 30 percent of the INCITE core-hours be run on eight or more racks (16.7 percent), and High Capability requires a minimum of 10 percent of the INCITE core-hours be run on 16 or more racks (33.3 percent). See Appendix A for more detail on the capability calculation. Table 2.11 and Figure 2.7 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 2.8 shows the overall distribution of job sizes over time, while Figure 2.9 shows the capability hours used over time by program.

Table 2.11 Capability Results for Mira

Mira (Blue Gene/Q) 48K-node, 768K-core, 1.6 GHz, 768 TB RAM						
	CY 2016			CY 2017		
Capability Usage	Total Hours	Capability Hours	Capability (%)	Total Hours	Capability Hours	Capability (%)
INCITE Overall	3.9B [a]	3.0B	76.4	4.5B [e]	3.2B	72.2
INCITE High	3.9B [a]	1.6B	40.9	4.5B [e]	1.2B	27.5
ALCC Overall	1.7B [b]	1.0B	60.0	1.6B [f]	1.0B	65.7
ALCC High	1.7B [b]	327.0M	19.2	1.6B [f]	262.7M	16.6
Director's Discretionary Overall	365.5M [c]	158.1M	43.3	309.6M [g]	155.3M	50.2
Director's Discretionary High	365.5M [c]	78.7M	21.5	309.6M [g]	65.9M	21.3
TOTAL Overall	6.0B [d]	4.2B	69.7	6.4B [h]	4.4B	69.5
TOTAL High	6.0B [d]	2.0B	33.6	6.4B [h]	1.6B	24.5

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

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

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

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

[e] Usage includes 3.0M core-hours from Cetus production jobs.

[f] Usage includes 2.7M core-hours from Cetus production jobs.

[g] Usage includes 7.7M core-hours from Cetus production jobs.

[h] Usage includes 13.4M core-hours from Cetus production jobs.

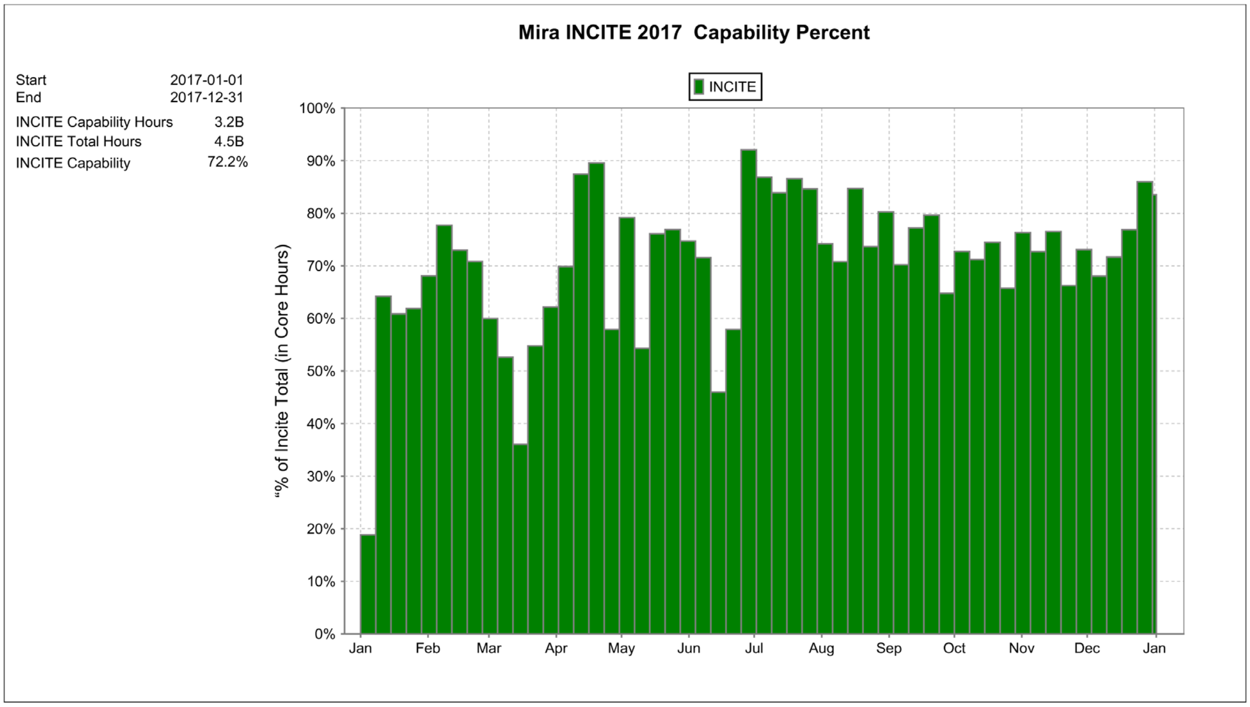

Figure 2.7 Mira's Overall Capability for INCITE

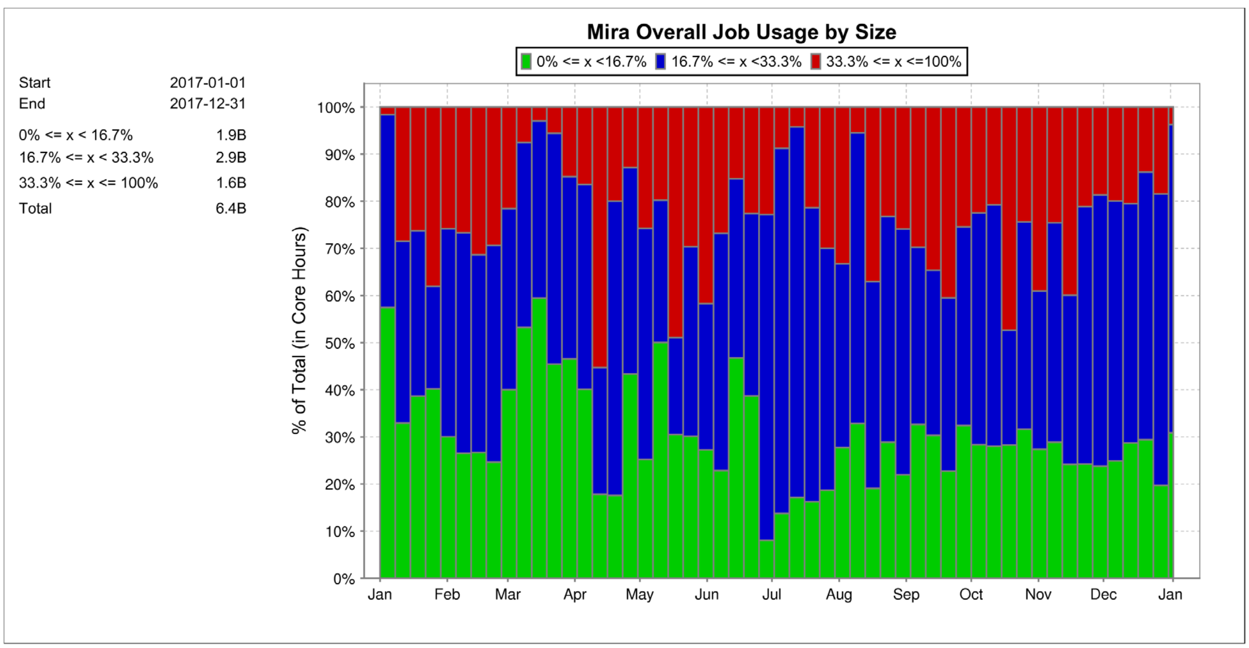

Figure 2.8 Mira's Job Usage by Size

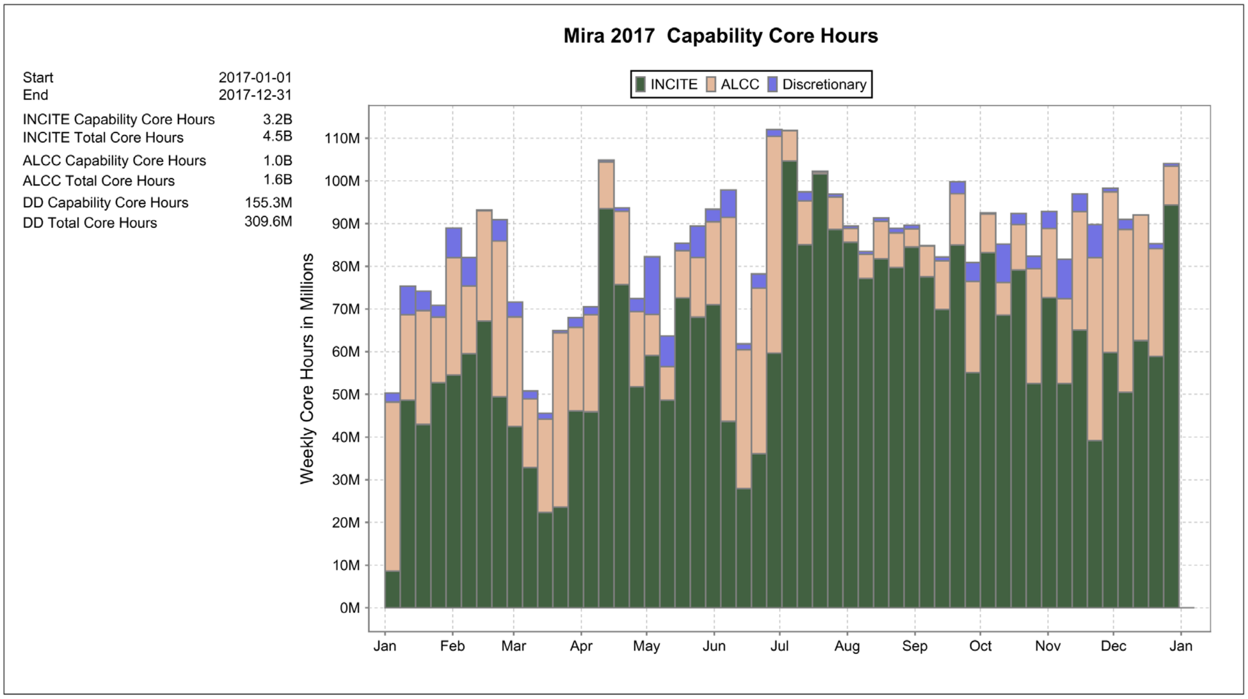

Figure 2.9 Mira's Capability Usage by Award Category

Conclusion

The ALCF is maximizing the use of its HPC systems and other resources consistent with its mission. We have 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 the National Energy Research Scientific Computing Center (NERSC), and the values reported are reasonable. These measures are summarized in Table 2.12.

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 2017, this regular failure analysis has continued to drive code changes to Cobalt, ALCF's job scheduler, and has provided details to support debugging of storage system problems, as noted in Section 2.1.1 under "Explanation of Significant Availability Losses" for both Theta and Mira.

Table 2.12 Summary of All Metrics Reported in the Business Results Section

	Mira (Blue Gene/Q): 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				Theta (Cray XC40): 3240-node, 21K-core, 56 TB MCDRAM, 679 TB DDR4 [h]	
	CY 2016		CY 2017		CY 2017	
	Target	**Actual**	**Target**	**Actual**	**Target**	**Actual**
Scheduled Availability	90.0%	98.9%	90.0%	99.5%	90.0%	98.6%
Overall Availability	90.0%	94.9%	90.0%	96.4%	80.0%	90.0%
System MTTI	N/A [c]	10.00 days	N/A	10.37 days	N/A	6.94 days
System MTTF	N/A	40.55 days	N/A	36.44 days	N/A	30.30 days
INCITE Usage	3.5B	3.9B [d]	3.5B	4.5B [f]	N/A	N/A
Total Usage	N/A	6.0B [e]	N/A	6.4B [g]	N/A	724.1M
System Utilization	N/A	90.8%	N/A	95.5%	N/A	82.7%
INCITE Overall Capability [a]	40.0%	76.4% [c]	40.0%	72.2% [f]	N/A	N/A
INCITE High Capability [b]	10.0%	40.9% [d]	10.0%	27.5% [f]	N/A	N/A

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

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

[c] N/A = Not applicable.

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

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

[f] Usage includes 3.0M core-hours from Cetus production jobs.

[g] Usage includes 13.4M core-hours from Cetus production jobs.

[h] See Appendix A for the definition of the Theta production node count calculation.

This page intentionally left blank.

Section 3. Strategic Engagement and Results

(a) Are the processes for engaging with programmatic and strategic stakeholders (i.e., beyond the user population) effective and do these processes enable scientific outputs and accomplishments consistent with the DOE strategic goals?

(b) Is the allocation of resources reasonable and effective?

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 an HPC environment and have enabled scientific achievements that would not have been possible otherwise.

In this section, ALCF reports:

- Science Output: users' peer-reviewed publications;
- Scientific Accomplishments;
- Allocation of Resources: Facility Director's Discretionary Reserve Time;
- Stakeholder Engagement and Outreach.

3.1 Scientific Output

The Facility tracks and reports annually the number of refereed publications resulting (at least in part) from use of 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.

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 one publication in *Nature*, one in *Nature Methods*, and one in *Nature Communications*. In addition, ALCF users published in other high-quality, peer-reviewed publications such as the journal *Science*, *Proceedings of the National Academy of Sciences (PNAS)*, *Physical Review Letters,* and the proceedings of the *International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*.

Table 3.1 Summary of Users' Peer-Reviewed Publications

Nature Journals	Science	PNAS	Physical Review Letters	SC17	Total 2017 Publications
3	2	2	4	4	211

3.2 Scientific Accomplishments

3.2.1 Molecular Design of Dye-Sensitized Solar Cells with Data Science

The Science

Dye-sensitized solar cells (DSSCs) are a next-generation photovoltaic technology whose transparent and low-cost nature make them a particularly strong contender for "smart windows" — windows that generate electricity from sunlight. Yet, like so many solar-cell technologies, innovation is being held up by a lack of suitable materials for these devices. This project marries the latest technical capabilities in natural language processing, machine learning, and quantum-chemical calculations to the world-leading supercomputing resources available at Argonne. The overarching concept is to search through a representative set of all possible chemical molecules and use artificial intelligence to target the chemicals whose molecules have optical properties that would yield optimum device function in DSSCs.

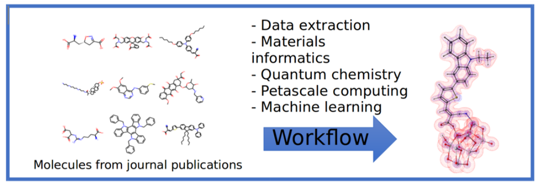

Figure 3.1 Schematic Diagram of Computational Workflow

(Image courtesy of Álvaro Vázquez-Mayagoitia, Argonne National Laboratory.)

The Impact

Buildings account for 40 percent of the total energy usage in the United States. "Smart windows" that generate electricity from sunlight hold exciting prospects for meeting entire cities' building energy demands in a fully sustainable fashion. DSSCs have a design point trade-off between cost and efficiency that can be explored through data science techniques and leadership computing.

Summary

Organic DSSCs are a low-cost alternative to traditional solar cells that can produce electricity and be functional as well. The combinatorial space of chemical building blocks to produce efficient dyes makes it unmanageable to experimentally study them. This project combines experimental data, molecular simulation, and data science techniques to reduce the search space, vastly reducing the overall time required to find the best dyes. Experimental data are obtained from hundreds of thousands of publications with data mining tools and stored in a database. An automated workflow pulls chemical formulas from the database of experimentally known dyes and produces 3D structures in order to simulate them with quantum chemistry (Figure 3.1). Following this, experimental and theoretical information is combined and used to inform machine learning models to predict efficient candidate dyes. This project will produce one of the largest datasets of experimental and theoretical information and will serve as a valuable resource for new dyes. The team collaborated with ALCF staff to develop the parallel version of the term data extraction tool and to accelerate the data extraction with optimized natural language processing tools. ALCF also developed a novel workflow software that provides physical chemical properties from quantum chemistry methods.

Contact

Jacqueline M. Cole
University of Cambridge and Argonne National Laboratory
jmc61@cam.ac.uk

Publication

Cole, J. M., *et al.* Discovery of S···C≡N Intramolecular Bonding in a Thiophenylcyanoacrylate-Based Dye: Realizing Charge Transfer Pathways and Dye···TiO_2 Anchoring Characteristics for Dye-Sensitized Solar Cells. *ACS Appl. Mater. Interfaces* **9** (31), 25952–25961 (2017). [DOI:10.1021/acsami.7b03522]

Related Links

ASCR Discovery research highlight: A window into solar [http://ascr-discovery.science.doe.gov/2017/09/a-window-into-solar/]

Highlight Categories

Performer/Facility: ASCR-ALCF, BES-APS

Date Submitted to ASCR

October 2017

3.2.2 Multiscale Simulations of Human Pathologies

The Science

Blood clots play a significant role in aortic aneurysms and dissections, but the mechanism is not fully understood. The first data-driven, multiscale, multiphysics model of the biomechanics of thrombus in aortic dissection is being developed to further investigate the role of blood clots.

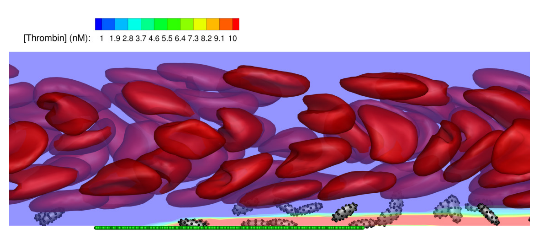

Figure 3.2 A snapshot of the dissipative particle dynamics simulation showing red blood cells, platelets, von Willebrand Factor ligands (green particles on the lower wall), and contours of thrombin concentration from the coagulation cascade. Also shown are platelets that have adhered to the injured area.

(Image courtesy of George Karniadakis and Alireza Yazdani, Brown University.)

The Impact

A better understanding of the roles of thrombus in aortic dissection could lead to improved testing and treatments for the disease. In addition, findings from this project will have important implications for a host of other vascular conditions, providing information that could contribute to improved treatments for a broad class of medical problems.

Summary

Thoracic aortic aneurysms leading to acute aortic dissections (TAADs) are a common cause of premature death in the United States. A Brown University team has developed the first data-driven, multiscale, multiphysics model of the biomechanics of blood clots, or thrombi, in aortic

dissections. The team has developed a flexible approach in which multiple computer codes are integrated to perform truly multiscale simulations of realistic biological systems. The team has also developed the numerical methods to address clot formation and platelet aggregation in complex geometries (Figure 3.2). This combination of fluid mechanics, biology and chemistry concepts, and corresponding simulation codes has allowed researchers to make significant progress in the modeling of red blood cells in healthy and diseased conditions. In addition, a multiscale framework has been developed to couple these heterogeneous continuum and atomistic solvers to address the challenging multiscale problems, especially those in pathological conditions such as TAAD. The team collaborated with ALCF scientists to optimize I/O and to develop a pipeline for concurrent simulation and visualization.

Contact

George E. Karniadakis
Brown University
george_karniadakis@brown.edu

Publications

Yazdani, A., and Karniadakis, G. E. Sub-cellular Modeling of Platelet Transport in Blood Flow through Microchannels with Constriction. *Soft Matter* **12**, 4339–4351 (2016). [DOI:10.1039/c6sm00154h]

Yazdani, A., *et al.* Dynamic and Rheological Properties of Soft Biological Cell Suspensions. *Rheologica acta* **55**, 433–449 (2016). [DOI:10.1007/s00397-015-0869-4]

Chang, H. Y., *et al.* MD/DPD Multiscale Framework for Predicting Morphology and Stresses of Red Blood Cells in Health and Disease. *PLoS Comput. Biol.* **12** (10), e1005173 (2016). [DOI:10.1371/journal.pcbi.1005173]

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR

June 2017

3.2.3 Direct Numerical Simulation of Compressible, Turbulent Flow

The Science

Numerical simulations are revealing the origin of the large-scale, low-frequency pressure fluctuations that lead to severe structural fatigue, shorter lifespan, and even catastrophic failure in high-speed aircraft. This INCITE project, involving direct numerical simulation of supersonic turbulent flow, aims to provide an understanding of this phenomenon.

The Impact

Most of us have experienced a long, tedious flight, and have dreamed of faster aircraft. While much progress has been made toward routine flight at higher speeds, the presence of intense, unsteady aerodynamic loads remains a fundamental design difficulty. With this new understanding of the flow physics, researchers are exploring flow modification and control to enable a new generation of fast, reliable aircraft.

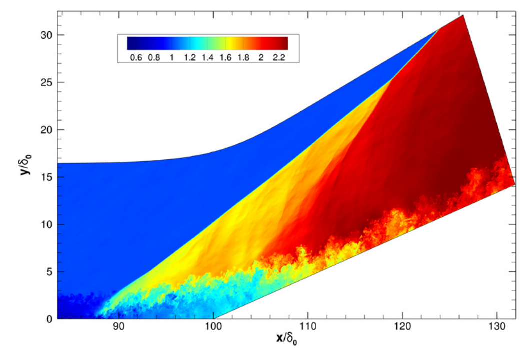

Figure 3.3 Instantaneous density field in a ramp flow at Mach 2.3 over a 24° ramp. An equilibrium turbulent boundary layer forms at x/δ_0 = 80, which is just off the left of the figure. This turbulent boundary layer provides a standard inflow condition for the separation that occurs between x/δ_0 = 80 and the corner at x/δ_0 = 100. The turbulent boundary layer is visible as the irregular region at the bottom of the image; it separates ahead of the corner and reattaches on the ramp downstream. The transition from low to high density away from the corner indicates the presence of the oblique shock wave. Acoustic waves are visible, radiating from the boundary layer turbulence and interacting with the shock. These waves can only be resolved on a fine mesh with a high-order numerical method.

(Image courtesy of Jonathan Poggie, Purdue University.)

Summary

Two conceptual models of separation unsteadiness have been advocated in the literature: an oscillator viewpoint and an amplifier viewpoint. In the oscillator model, the separated flow has an inherent instability, which leads to self-excited oscillations. In the amplifier model, disturbances in the incoming flow are selectively amplified in the separated flow. The separation region has a characteristic space and time scale, and it reacts strongly to disturbances on the order of these scales, but not to smaller-scale disturbances. This work aims to settle the long-standing question of amplifier versus oscillator. Because of the large range of scales between fine-grained turbulence and separation bubble motion, this research requires a few large calculations that will run for a very long time. To this end, this team has developed an adaptive, compact-difference code that can capture both fine-grained turbulence and shock waves (Figure 3.3). Multi-level parallelism in the code structure enables very fast execution on Mira. The team has carried out the first-ever fully resolved computation of a shock-wave/boundary-layer interaction, which enabled them to study the interaction of the shock with turbulence-generated sound. Computational results reproduced the spectrum predicted by the Plotkin (1975) amplifier theory of separation unsteadiness. The project team is now studying an analogous case with a laminar inflow and searching for evidence of large-scale separation bubble motion in the absence of incoming disturbances. They are also exploring flow control, that is, the effects on the separation

bubble of controlled disturbances in the incoming flow. ALCF staff helped with porting Fortran 2008 code from Intel's ifort to IBM's xlf compiler, changing from the Lustre file system to IBM GPFS, and with plans for migration to new hardware.

Contact

Jonathan Poggie
Purdue University
jpoggie@purdue.edu

Publications

Porter, K. M., and Poggie, J. Turbulence Structure and Large-Scale Unsteadiness in Shock-Wave/Boundary Layer Interaction. 55th AIAA Aerospace Sciences Meeting, AIAA SciTech Forum (AIAA 2017-0533) (2017). [https://doi.org/10.2514/6.2017-0533]

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR

January 2018

3.2.4 Simulating Macromolecular Processes under Realistic Conditions through High-Performance Computing

The Science

PH is an acidity/alkalinity scale and one of the primary characteristics of cellular environments. PH plays a significant role in macromolecular interactions and also plays an important part in drug optimization.

The Impact

Results from this project yield insight into the origin of acidity/alkalinity effects on drug adsorption. The scalable and generalizable feature of the constant-pH method opens the door for pH-dependent drug optimization and discovery.

Summary

Computational biophysicists from the University of Chicago and Argonne National Laboratory implemented state-of-the-art algorithms for constant-pH molecular dynamics simulations. The implementation is massively parallel and works efficiently for large biomolecules under realistic pH conditions. The team explored the structural response of a small drug molecule under various acidity/alkalinity conditions and revealed how pH conditions affect translocation of small molecules across cell membranes (Figure 3.4). In collaboration with the University of Illinois at Urbana-Champaign and UChicago, Early Science postdoc Brian Radak, supervised by staff member Wei Jiang, developed and implemented a scalable algorithm for constant-pH simulations in NAMD. The relevant code was officially released in NAMD version 2.12. Staff and postdocs collaborated on the INCITE proposal.

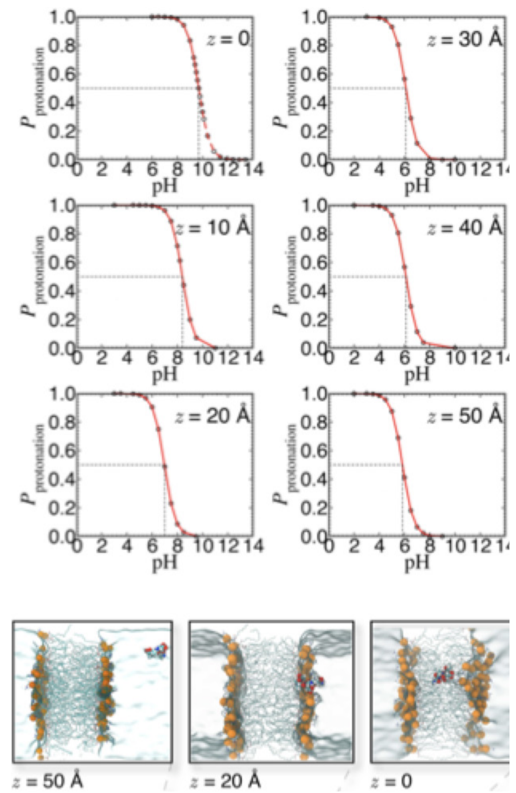

Figure 3.4 The translocation of a titratable mini-peptide across a cell membrane was simulated at various separations (lower three panels) and constant pH. The insets depict the protonation probability of the central residue for different positions along the bilayer normal (upper six panels). It can be clearly seen how the response of the protonated state to the pH conditions changes across the membrane. The variation in protonation is accompanied by significant conformational changes in the biomolecule.

(Image courtesy of Brian Radak, Argonne National Laboratory.)

Contact

Benoît Roux
The University of Chicago and Argonne National Laboratory
roux@uchicago.edu

Publications

Radak, B., *et al.* Constant-pH Molecular Dynamics for Large Biomolecular Systems, *J. Chem. Theory and Comput.* **13**, 12 (2017). [DOI: 10.1021/acs.jctc.7b00875]

Radak, B., *et al.* Inclusion of ph Effects in Molecular Dynamics Simulations of Membranes and Membrane Proteins. *Biophys J.* **110**, 646a (2016). [DOI: 10.1016/j.bpj.2015.11.3456]

Related Links

Argonne National Laboratory research highlight: Theta provides insights into the building blocks of life [https://www.alcf.anl.gov/articles/theta-provides-insights-building-blocks-life]

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR

March 2018

3.2.5 Atomistic Simulations of Nanoscale Oxides and Oxide Interfaces

The Science

Corrosion, the process that leads to rust, is a particular kind of metal oxidation reaction. If it can be carefully controlled at the atomic scale, it could offer scientists ways to advance state-of-the-art battery or drug delivery technologies.

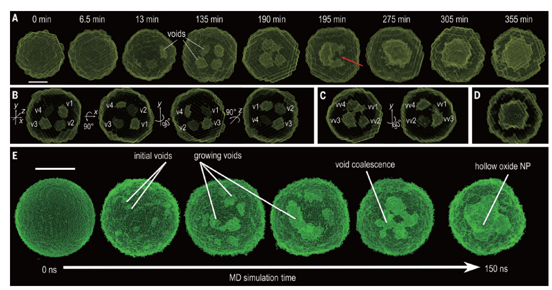

Figure 3.5 (A–D) Structures reconstructed from SAXS patterns. (E) Snapshots of the 3D structure of iron nanoparticles in the course of the oxidation process, captured through large-scale reactive molecular dynamic simulations.

(Image courtesy of Yugang Sun, Temple University; Xiaobing Zuo, Subramanian Sankaranarayanan, Sheng Peng, Badri Narayanan, and Ganesh Kamath, Argonne National Laboratory.)

The Impact

Multimillion-atom simulations on Mira provided scientists with an atom-scale "movie" of the oxidation process that transforms iron nanoparticles into nanoshells. This fundamental chemical reaction is relevant to study of advanced batteries and drug-delivery technologies.

Summary

A research team studied the behavior of iron nanoparticles by watching them in real time as they oxidized. By using a combination of X-ray scattering and computational simulation, the team was able to observe and model the changes in nanoparticle geometry as they occurred. This knowledge adds to scientists' understanding of fundamental reactions like corrosion and helps in building a foundation for developing new techniques that researchers can use to image, manipulate, and control such reactions at the atomic scale. In their study, which was published in the journal *Science*, the team sought to understand the Kirkendall effect in small particles of iron during oxidation at the nanoscale level. At this scale, iron nanoparticles exposed to an oxygen environment exhibit a unique property — they form exotic structures, such as hollowed-out nanoparticles called nanoshells (Figure 3.5), which already have been used as electrodes in battery applications and as vehicles for drug delivery in medicine. The shape, structure, and distribution of the holes in these nanoshells depend on how oxidation progresses in time. The X-rays helped the scientists monitor structural evolution as the nanoparticles oxidized. By using small- and wide-angle X-ray scattering (SAXS, WAXS), the researchers could see how voids form at a relatively high spatial resolution; the resolution nevertheless did not reach to the level of tracking individual atoms. For this insight, the team turned to the ALCF. LAMMPS simulations provided more detailed insight into the transformation of nanoparticles into nanoshells and the atomic-scale processes that govern their evolution. ALCF staff collaborated with IBM, Lawrence Berkeley National Laboratory, and Sandia National Laboratories to optimize LAMMPS.

Contact

Subramanian Sankaranarayanan
Argonne National Laboratory
ssankaranarayanan@anl.gov

Publications

Sun, Y., *et al.* Quantitative 3D Evolution of Colloidal Nanoparticle Oxidation in Solution. *Science* **356**, 6335, 303–307 (2017). [DOI: 10.1126/science.aaf6792]

Aktulga, H. M., *et al.* Optimizing the Performance of Reactive Molecular Dynamics Simulations for Multi-Core Architectures. *Int. J. High Perform. C.* (in press). [https://arxiv.org/abs/1706.07772]

Related Links

Argonne National Laboratory research highlight: New study reveals the mystery behind the formation of hollowed nanoparticles during metal oxidation [http://www.anl.gov/articles/new-study-reveals-mystery-behind-formation-hollowed-nanoparticles-during-metal-oxidation]

Highlight Categories

Performer/Facility: ASCR-ALCF, BES-APS

Date Submitted to ASCR

July 2017

3.2.6 Predictive Modeling of Functional Nanoporous Materials, Nanoparticle Assembly, and Reactive Systems

The Science

Micropore channels in membranes allow certain molecules to pass through them but not others. Critical work in this area is focused on developing efficient synthetic procedures to grow thin molecule-selective membranes with tunable micropore dimensions.

The Impact

Membrane-based separations provide a route to purifying mixtures at a fraction of the energy cost compared to traditional distillation processes, which would have a significant impact on lowering global energy use, emissions, and pollution. Simulations on Mira supported and validated experimental findings, which identified a simpler and more cost-effective synthesis route for these ultra-selective high-flux membranes.

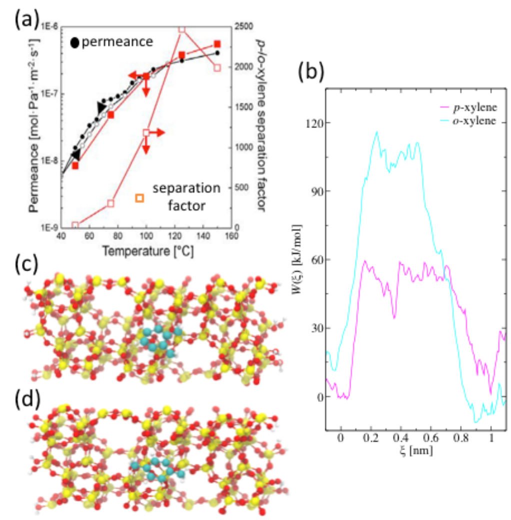

Figure 3.6 (a) Experimental PX permeance and selectivity over o-xylene (OX) for MFI nanosheet membrane, (b) PMFs for p-xylene (PX) and OX along the straight channel, and snapshots of (c) OX and (d) PX near the top of their barriers.

(Image courtesy of Evgenii Fetisov and J. Ilja Siepmann, University of Minnesota.)

Summary

MFI-type zeolites are aluminosilicate (or silicate) materials with three-dimensionally-connected pores of varying width — giving rise to molecular recognition features — and are widely used in the chemical industry as selective catalysts or adsorbents. While zeolite materials (e.g., MFI) have been shown to efficiently and selectively separate mixtures, their application as membranes is plagued by challenges to synthesize membranes of sufficient size and quality. The usual procedure to generate single MFI nanosheets is a costly, inefficient multi-step exfoliation-centrifugation process where individual layers of material are peeled away until a single sheet remains. In this work, the team developed a procedure to directly synthesize individual nanosheets using nanocrystal seeds to grow individual MFI sheets with a thickness of 5 nanometers, sidelengths of 1–2 micrometers, and preferred orientation with straight micropore channels connecting both flat surfaces. The team used several experimental techniques to characterize the structure of these high aspect ratio nanosheets, including X-ray diffraction measurements collected at the Advanced Photon Source, as well as the large-scale computing resource Mira to assist with interpreting experiments and predicting membrane performance. The newly synthesized membranes yielded separation factors of about 2500 at 125°C for mixtures of o-xylene (OX) and p-xylene (PX) (Figure 3.6), which differ only in the relative placement of two methyl groups around a benzene ring. The highest literature value previously reported was 1000 (i.e., a 2.5× improvement). Theoretical calculations corroborated proposed mechanisms for the kinetically controlled growth of these nanosheets as opposed to strain-induced processes. In addition, Mira was used to compute free energy barriers for molecule transport through the nanosheet for both OX and PX molecules using quantum mechanical simulations, which supported experimental measurements and predicted that even higher separations can be achieved. ALCF staff member Chris Knight assisted the team with developing ensemble submission scripts to enhance the workflow of capability-scale calculations (131,072 processors, or greater) on Mira.

Contact

J. Ilja Siepmann
University of Minnesota
siepmann@umn.edu

Publications

Jeon, M. Y., *et al.* Ultra-selective High-flux Membranes from Directly Synthesized Zeolite Nanosheets. *Nature* **543**, 690–694 (March 30, 2017). [DOI:10.1038/nature21421]

Related Links

University of Minnesota research highlight: Researchers develop groundbreaking process for creating ultra-selective separation membranes [https://cse.umn.edu/news-release/researchers-develop-groundbreaking-process-creating-ultra-selective-separation-membranes/]

Highlight Categories

Performer/Facility: ASCR-ALCF, BES-APS

Date Submitted to ASCR

May 2017

3.2.7 Kinetic Simulations of Relativistic Radiative Magnetic Reconnection

The Science

Intense X-ray and gamma-ray flares come from the extreme environments surrounding pulsars and supermassive black holes. The processes generating these intense radiation events are not well understood. This work makes significant progress toward understanding how this radiation is produced.

The Impact

Given the two viable mechanisms for producing the aforementioned intense radiation, this work has used state-of-the-art large-scale simulations to demonstrate that both magnetic reconnection and turbulent acceleration can contribute to the observed radiation. This furthers our understanding of these far-away objects, helping us to use their properties to understand larger properties of the universe, and also demonstrates our ability to accurately model highly energetic plasmas.

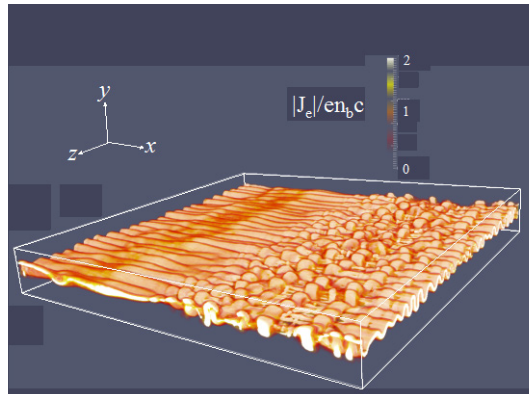

Figure 3.7 The current sheet in magnetic reconnection becomes distorted, rippling in the y-z plane, an effect that cannot be captured by two-dimensional simulations. Despite the distortion, particles in the current sheet experience significant acceleration owing to magnetic reconnection. (Image appears in ApJL.)

(Image courtesy of Greg Werner, University of Colorado Boulder.)

Summary

Intense X-ray and gamma-ray flares have been observed emanating from extreme environments surrounding distant astrophysical objects such as pulsars and supermassive black holes. This high-energy radiation is believed to originate from electrons — and sometimes positrons — that have been accelerated to extremely high, ultrarelativistic energies (perhaps thousands of times higher than in the largest man-made particle accelerators). Two leading candidates for acceleration mechanisms are magnetic reconnection and turbulence. Magnetic reconnection converts magnetic energy into particle kinetic energy through a specific rearrangement of the magnetic field, disconnecting magnetic field lines and “reconnecting” them in a different, lower-energy configuration. Turbulence, on the other hand, dissipates magnetic and fluid-flow energy in charged-particle plasmas via a complex multi-scale cascade, which transfers energy from large-scale flows to microscopic dissipation processes. Reconnection and turbulence, despite fundamental differences, may in fact be related: reconnection can generate turbulence, while turbulence may dissipate magnetic energy through magnetic reconnection. Although their specific motivation is astrophysical, similar mechanisms are thought to power solar flares and related X-ray emissions from our sun. With a campaign of very large, computationally intensive 3D simulations, this work has demonstrated systematically that, despite potentially disruptive 3D instabilities, particle acceleration is nearly the same in 3D reconnection as in 2D approximations (Figure 3.7). In other words, magnetic reconnection is a viable mechanism for astrophysical particle acceleration. In addition, using large-scale 3D particle-in-cell simulations of driven turbulence in relativistic electron-positron plasmas, the researchers have demonstrated that turbulence can indeed be an efficient particle accelerator. Turbulence in this high-energy, relativistic regime has proven to be highly amenable to massively parallel computation, promising to yield valuable developments in future studies of kinetic processes in turbulence. This project team worked directly with ALCF's catalyst, visualization, and performance engineering teams.

Contact

Dmitri Uzdensky
University of Colorado Boulder
uzdensky@colorado.edu

Publications

Zhdankin, V., *et al.* Kinetic Turbulence in Relativistic Plasma: From Thermal Bath to Nonthermal Continuum. *Phys. Rev. Lett.* **118**, 055103 (2017). [DOI:10.1103/PhysRevLett.118.055103]

Werner, G. R., *et al.* Nonthermal Particle Acceleration in Collisionless Relativistic Electron-proton Reconnection. *Mon. Not. R. Astron. Soc.* **473**, 4840-4861 (2018). [DOI:10.1093/mnras/stx2530].

Werner, G. R., and Uzdensky, D. A. Nonthermal Particle Acceleration in 3D Relativistic Magnetic Reconnection in Pair Plasma. *Astrophys. J. Lett.* **843**, L27 (2017). [DOI:10.3847/2041-8213/aa7892]

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR

August 2017

3.3 Allocation of Resources

3.3.1 Usage of the INCITE and ALCC Hours

The INCITE 2017 program allocated 3.5 billion core-hours on Mira. The allocation usage is shown in Figure 3.8. Of the 30 INCITE projects, 20 projects used more than 90 percent of their allocation. Seventeen projects used their entire allocation (or more), including seven projects using more than 150 percent, of which three projects used more than 200 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, 2017.

Of the remaining 10 projects with outstanding allocations, six projects used more than 50 percent of their time and of those six, three projects used more than 70 percent of their time. Only four projects used less than 50 percent of their allocation. A total of 4.5 billion core-hours was delivered to INCITE almost entirely on Mira. The contribution from Cetus to the total INCITE hours delivered was 0.07 percent.

For the 2016–2017 ALCC year, 29 projects had allocations on Mira for a total of 1.5 billion core-hours. The allocation usage is shown in Figure 3.9. Fifteen of these projects used 90 percent or more of their allocations, including one project that also used Cetus for production runs that were not easily accommodated on Mira. Cetus usage accounted for about 5.6 percent of that team's usage and less than 0.1 percent of the total ALCC 2016–2017 usage. Of the remaining 14 projects, three used 75 percent of their allocation or more, and only eight projects used less than 50 percent of their allocations.

The 2017–2018 ALCC year is approximately halfway through its allocation cycle. So far, 32 projects have received allocations of 1.1 billion core-hours across Mira and Theta. The projects have used a total of 656 million core-hours from July 1, 2017 through December 31, 2017. The allocation usage is shown in Figures 3.10 and 3.11.

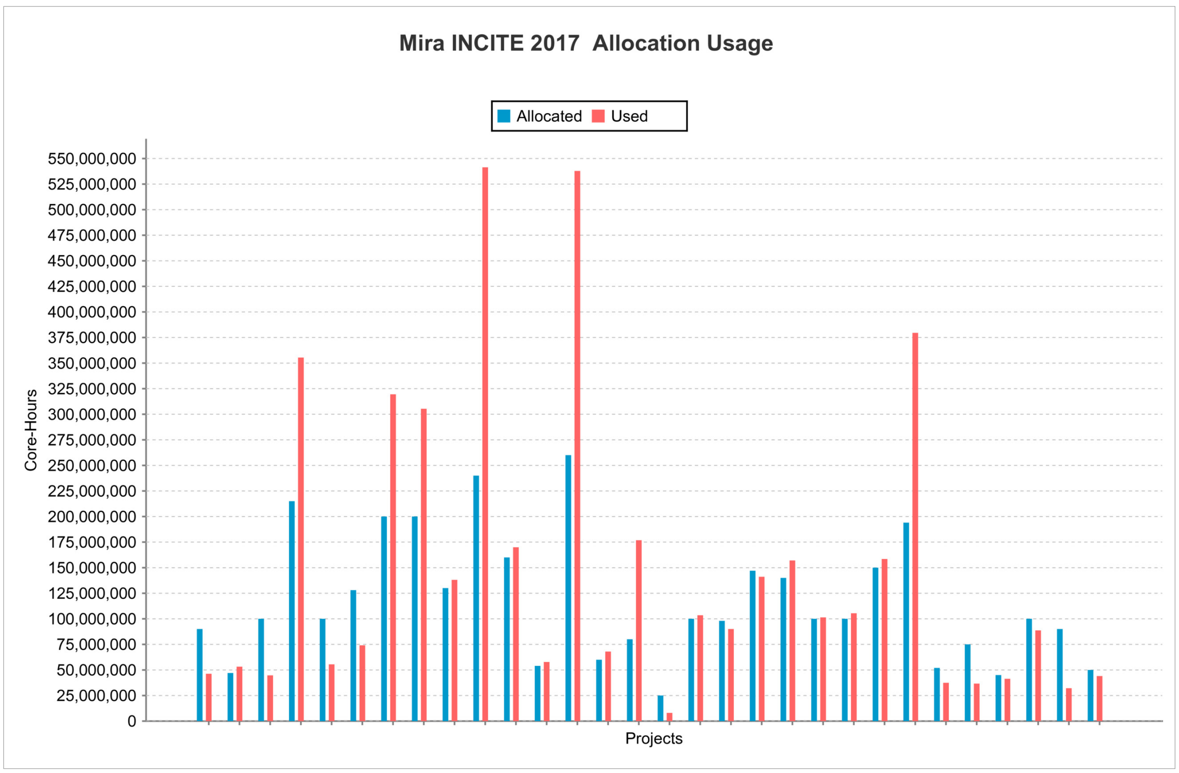

Figure 3.8 Mira INCITE 2017 Allocation Usage

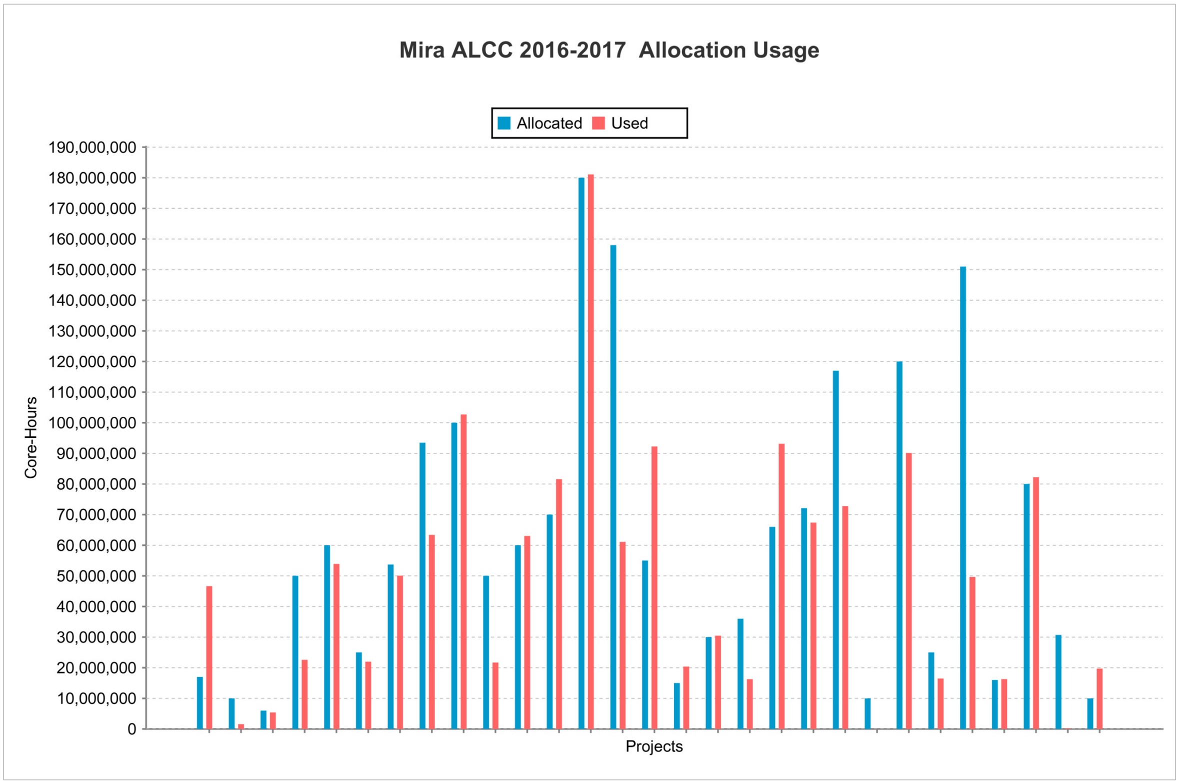

Figure 3.9 Mira ALCC 2016–2017 Allocation Usage

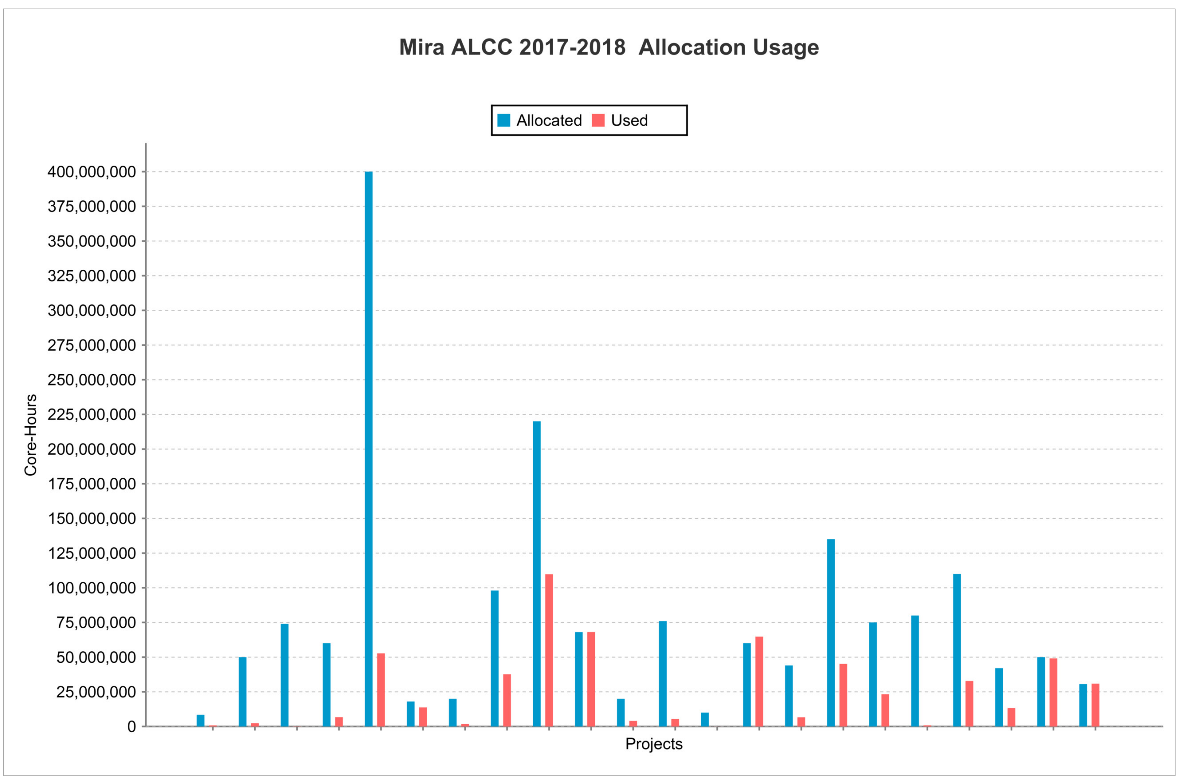

Figure 3.10 Mira ALCC 2017–2018 Allocation Usage

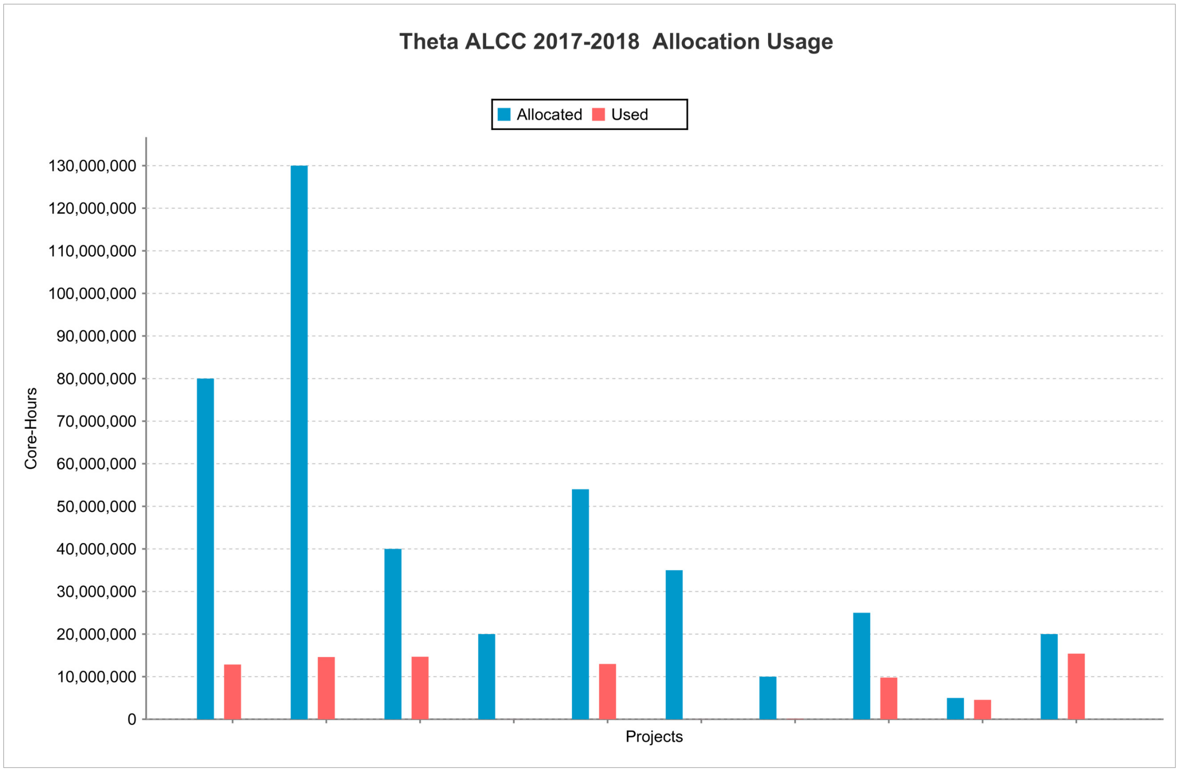

Figure 3.11 Theta ALCC 2017–2018 Allocation Usage

3.3.2 Facility Director's Discretionary Reserve Time

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 the 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 and Theta during 2017. 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, an approach that is in line with the ALCF core mission.

Table 3.2 DD Time Allocated and Used on Mira and Theta, 2017

Projects	Mira	Theta
Number of Projects	146	123
Allocated Core-Hours	1.7B	1.2B
Used Core-Hours	309.6M [a]	639.2M

[a] Usage includes 8 M core-hours from Cetus production jobs.

A list of the CY 2017 DD projects, including title, PI, institution, and hours allocated, is provided in Appendices B and C.

Figures 3.12 and 3.13 provide a breakdown of the CY 2017 DD allocations by domain for Mira and Theta, respectively.

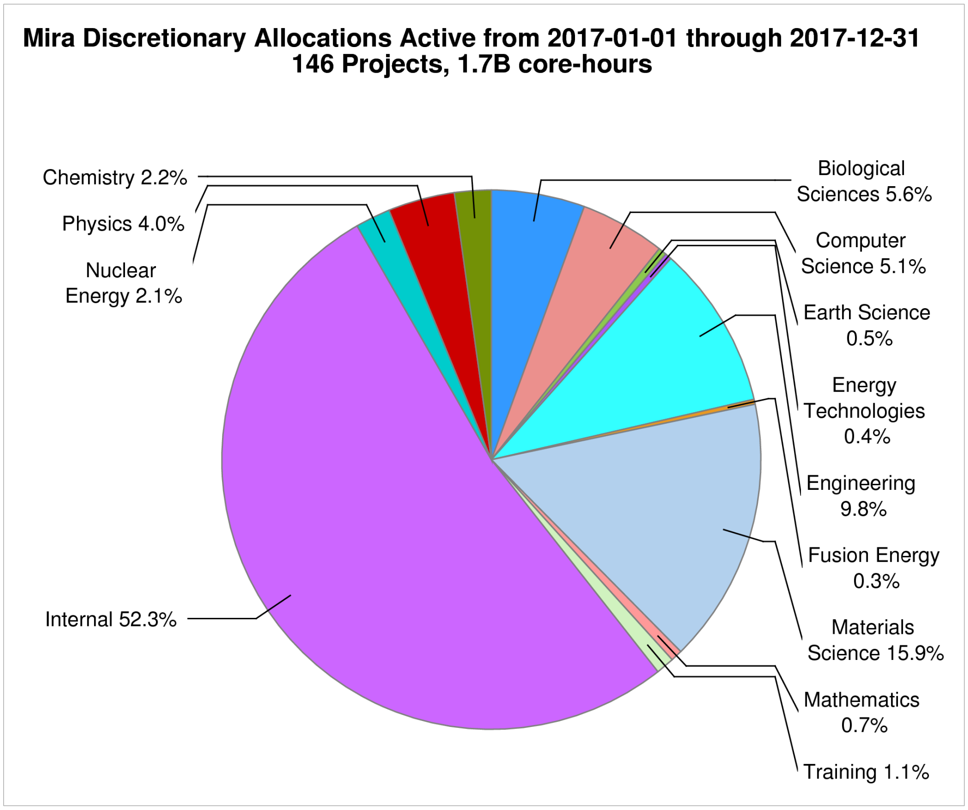

Figure 3.12 Mira CY 2017 DD Allocations by Domain

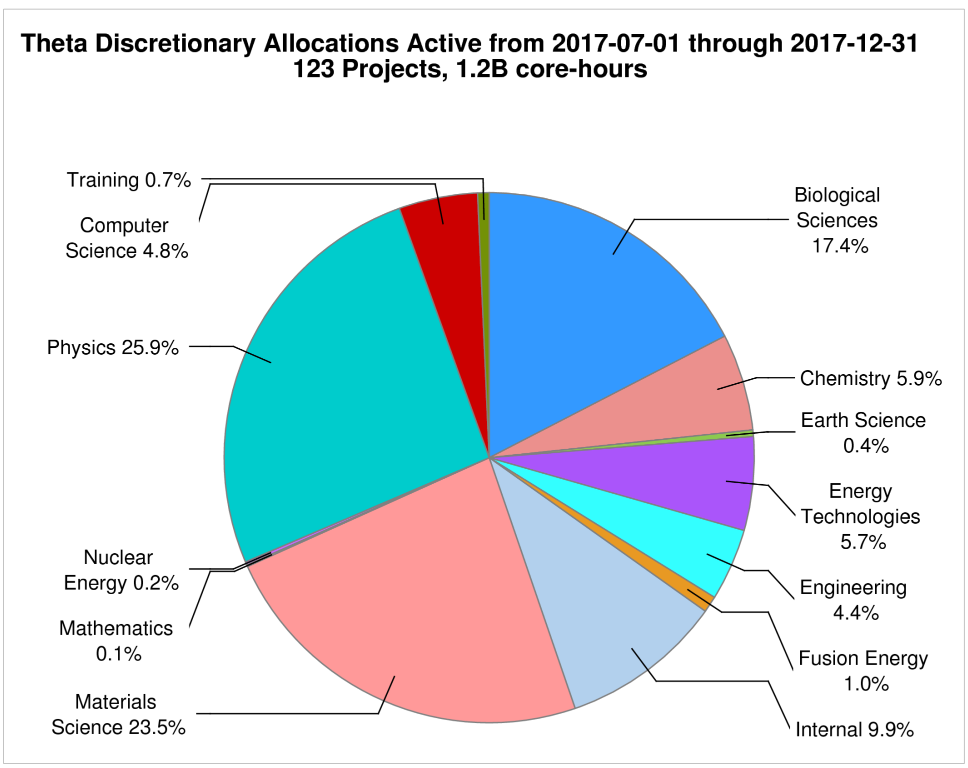

Figure 3.13 Theta CY 2017 DD Allocations by Domain

3.4 Stakeholder Engagement and Outreach

To help ensure that the ALCF delivers on its mission of breakthrough science at the largest computational scale, staff outreach needs tight engagement with domain science and a close eye on directions for supercomputing technologies. The ALCF provides a crucial balance of understanding how production science applications can move into new and exciting near, and future, architectures.

In 2017, the ALCF hosted multiple roadmap and deep-dive briefings from major and emerging vendors in addition to vendor discussions stemming from the CORAL-2 Request for Proposal (RFP) process and PathForward. The information from these briefings is used to drive acquisition decisions for Argonne's Joint Laboratory for System Evaluation (JLSE) and for ALCF projects that are relevant for both facility and application readiness. (JLSE activities are detailed in Section 4.2.1 of this report.)

ALCF outreach and engagement efforts with the science community are described in Section 3 and Section 4 of this report. ALCF staff participate in tight collaboration with both computer science and domains science projects to advance their use of ALCF resources — this happens for both production resources and for future resources through the Early Science Program. In addition to collaboration, staff members participate in community and domain activities. For example, ALCF staff led sessions at the 2017 March Meeting of the American Physical Society, the 2017 Meeting of the American Chemical Society, the Platform for Advanced Scientific Computing (PASC) Conference, and SC17. Staff conduct many university visits and town hall conversations each year to discuss areas of need with researchers. These activities focus not just on sharing ALCF work but collecting requirements and building collaborations.

ALCF staff are regular participants in DOE and National Science Foundation (NSF) workshops and reviews. Staff are engaged in standards committees and boards for both future and current technologies. In 2017, ALCF staff organized and ran the Fourth Workshop on the LLVM Compiler Infrastructure in HPC at SC.

ALCF staff has been closely participating with the Exascale Computing Project (ECP) since the start of the project. Over the course of 2017, the ALCF participated in multiple planning meetings with ECP and the other computing facilities (NERSC, OLCF) to develop the ECP/Facilities engagement plan. The ALCF sent 10 staff members to the ECP First Annual Meeting in addition to those attending as part of the ECP project. Staff participated in technical conversations on many projects and also engaged in specific Facility/ECP breakouts to discuss potential requirements for collaboration. In total, 28 ALCF staff have roles in multiple ECP projects and efforts including:

- Application Development Projects
- Software Technology Projects
- PathForward
- Application Development Working Group
- Software Technology Working Group
- Continuous Integration Working Group
- Industry Council

3.4.1 Outreach

Industry Lunch at SC Conference

The ALCF organizes an annual 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 members of ALCF's senior leadership to meet with industry leaders and discuss ways to work together better. Industry meeting attendees this year were a mixture of former users (Procter & Gamble, Boeing), current users (United Technologies, General Motors [GM], General Electric [GE]), and future users (Arcutek, British Petroleum [BP]). The meeting included an extended discussion on workforce development, wherein many of the industry participants encouraged the ALCF to increase training opportunities, especially around data analysis and machine learning.

Intel Xeon Phi User Group

David Martin, ALCF's industry partnerships and outreach manager, also serves as president of the Intel Xeon Phi User Group (IXPUG). IXPUG is an independent user group run by a steering group with members from major Intel Xeon Phi installations. Under David's leadership, IXPUG held eight significant meetings in 2017, including the spring and fall meetings. In addition, IXPUG held workshops and birds-of-a-feather sessions at the International Supercomputing Conference (ISC) and at SC. Regional and topic-specific meetings rounded out the year's activities.

Chicagoland HPC

Director Michael E. Papka and members of ALCF's science team are working with the University of Chicago's Research Computing Center, and others, to establish a Chicago-area framework for collaborations in HPC. The ALCF held a "Collaborations in Chicagoland HPC" kickoff workshop on October 4, 2017, and is in the process of setting up a website.

American Chemical Society (ACS) Symposium

ALCF staff member Yuri Alexeev and several other computational chemistry scientists from national laboratories, the U.S. Department of Defense, and the U.S. Army Corps of Engineers co-organized a petascale science session, "Extending Accuracy & Scales with Emerging Computing Architectures & Algorithms," as part of the Division of Computers in Chemistry (COMP) track at the 254th American Chemical Society National Meeting & Exposition in Washington, D.C., last August. The session was organized around the theme of novel algorithms for emerging computing architectures and featured talks by chemists, biologists, and material scientists who use HPC to solve specific, complex problems. The *International Journal of Quantum Chemistry* is publishing a special issue, guest edited by Yuri Alexeev and Robert Harrison, dedicated to HPC quantum chemistry methods presented during the session.

Resources for Performance Portability

ALCF staff member Tim Williams and other leads at NERSC and OLCF jointly created and deployed www.performanceportability.org, a website for resources and best practices about performance portability across the ASCR facilities. Tim and others led a breakout at the 2017 DOE Centers of Excellence Performance Portability Meeting, August 22–24, 2017, in Denver, Colorado, to demo and get feedback on the site.

Summer Student Research Programs

Every summer, ALCF staff members mentor undergraduate and graduate students on real-world research projects through DOE's Science Undergraduate Laboratory Internship (SULI) program and Argonne's Research Aide Appointments. In 2017, 39 students worked on projects ranging from system administration and data analytics to computational science and performance engineering.

Scaling to Petascale Institute

On June 26–30, 2017, the ALCF, National Center for Supercomputing Applications (NCSA)/Blue Waters Project, NERSC, OLCF, Stony Brook University, and the Texas Advanced Computing Center (TACC) organized a weeklong "Scaling to Petascale" Institute. The goal was to prepare participants to scale simulations and data analytics programs to petascale-class computing systems in support of computational and data-enabled discovery for all fields of study. Organizers worked to facilitate the participation of a large national audience. This feat was accomplished by streaming the sessions to a number of collaborating organizations using full-duplex audio/video connections. Sessions were also webcasted on YouTube Live. Recordings of the presentations were made available afterward. More than 400 people participated via 23 sites across the United States and in three other countries, and more than 200 people registered to watch the sessions via YouTube Live. The ALCF hosted nearly 40 attendees during the week and an additional 35 attendees on Monday, June 26, for the MPI Tutorial that was added as a public event to the institute this year. ALCF staff member Marta García was responsible for organization of the institute, as well as logistics and support of the ALCF host site.

Platform for Advanced Scientific Computing (PASC) Conference

ALCF staff member Ramesh Balakrishnan organized a mini-symposium, "High Order Discretizations in Large Eddy Simulation of High Reynolds Number Turbulent Flow in Engineering Applications," as part of the Engineering track at the 2017 Platform for Advanced Scientific Computing (PASC) Conference in Lugano, Switzerland, last June. The focus of this event was higher-order numerical discretizations, their impacts on the resolvable turbulent flow physics, and the scaling and parallel performance of higher-order discretizations on emerging computing hardware.

American Physical Society (APS) March Meeting

ALCF staff member Nichols Romero, Jack Wells (Oak Ridge National Laboratory), Jack Deslippe (NERSC), and Barry Schneider (National Institute of Standards and Technology) co-organized a petascale science session, "Petascale Science and Beyond: Applications and Opportunities for Materials Science," as part of the Division of Computational Physics (DCOMP) track at the 2017 March Meeting of the American Physical Society in New Orleans, Louisiana. The session was organized around this year's March Meeting theme of condensed matter physics, and also computational biology. Featuring three invited talks and 36 contributed talks, the petascale session was attended by researchers from universities, national laboratories, and industry. The contributed talks were presented by users of ALCF and other ASCR-funded computing facilities and also of NSF eXtreme Science and Engineering Discovery (XSEDE) resources.

Building New Academic Bridges

ALCF staff members Ira Goldberg and Beth Cerny sponsored a one-year student consulting project that provided sentiment analysis for ALCF's 2016 ALCF User Survey, as well as a flexible abridged survey. Graduate students from the Management Information Systems (MIS) consulting program at the University of Arizona's Eller College of Management provided frameworks and analysis that will be useful for future user surveys. For ALCF, it was an opportunity to reach out to a top-tier program outside of the Midwest.

3.4.2 Summary of Engagements with the Exascale Computing Project

The ALCF and the ECP have had many opportunities to engage in 2017, starting with the ECP's first annual meeting, January 29–February 3, 2017, in Knoxville, Tennessee, where ALCF staff members participated in the exascale systems and facilities meeting, technical talks, the poster session, and the industry council dinner.

Susan Coghlan, project director for future ALCF systems and David Martin, ALCF's manager of industry partnerships and outreach are both members of ECP's leadership team: Coghlan is deputy director of Hardware and Integration, and Martin is co-executive director of the Industry Council. ALCF staff members also participated in various 2017 working groups, including Katherine Riley (application development), Mark Fahey and Jini Ramprakash (facilities), Scott Parker (software technology), and Ben Lenard (continuous integration [CI]). Lenard, who served on both the general and technical working group for CI, also published his recommendations for building additional capabilities into CI solutions as a white paper.

The ECP was awarded a 2017–2018 ALCC allocation, which was used to support 24 small projects, or sub-allocations. ALCF staff provided initial on-boarding assistance for all the projects and ongoing support for any modifications requested by the ECP. ALCF staff also set up additional reporting to collaborate with ECP in gathering data about the individual projects and sent weekly e-mails to the project PIs. In addition, ALCF deployed a web page with information about the ECP projects, allocations, and usage by resource. A comma-separated values (CSV) file of all ECP jobs is also available at this site. A snapshot of the webpage is provided in Figure 3.14.

ECP ALCC 2017-2018 (subject to revision)

Generated on 2018-04-05

Summary from 2017-07-01 through 2017-12-31

			THETA					MIRA				
Project Name	PI	Field	Allocation	Hours Used	% Used	Total Capability (Hours)	Capability %(total)	Allocation	Hours Used	% Used	Total Capability (Hours)	Capability %(total)
CSC249ADCD01	Ian Foster	Computer Science	1,000,000	457,798	45.7%	0	0.0%	0	0	0.0%	0	0.0%
CSC249ADCD04	Tzanio Kolev, Mi Su..	Computer Science	0	0	0.0%	0	0.0%	4,000,000	1,232,678	30.8%	24,722	2.0%
CSC249ADCD05	Mahantesh Halappa..	Computer Science	3,000,000	1,730,721	57.6%	1,022,817	59.0%	0	0	0.0%	0	0.0%
CSC249ADCD502	Kenneth Roche	Computer Science	2,000,000	265,253	13.2%	0	0.0%	4,000,000	1,091,734	27.2%	945,284	86.5%
CSC249ADOA01	Rick L. Stevens	Biological Sciences	50,000,000	12,266,983	24.5%	2,357,812	19.2%	0	0	0.0%	0	0.0%
CSC249ADSE01	Salman Habib	Physics	7,680,000	2,195	0.0%	0	0.0%	0	0	0.0%	0	0.0%
CSC249ADSE03	Paul Mackenzie	Physics	16,500,000	6,708,570	40.6%	0	0.0%	20,000,000	6,954,221	34.7%	0	0.0%
CSC249ADSE04	Arthur F. Voter	Nuclear Energy	1,600,000	0	0.0%	0	0.0%	3,000,000	1,529	0.0%	0	0.0%
CSC249ADSE06	Jean-Luc Vay	Physics	4,000,000	288,564	7.2%	204,582	70.8%	0	0	0.0%	0	0.0%
CSC249ADSE08	Tom Evans	Nuclear Energy	20,000,000	845,724	4.2%	0	0.0%	200,000,000	107,389,617	53.6%	49,899,110	46.4%
CSC249ADSE09	Paul Kent	Materials Science	16,000,000	12,367,373	77.2%	3,934,674	31.8%	30,000,000	22,193,036	73.9%	20,876,092	94.0%
CSC249ADSE12	Amitava Bhattacharj..	Computer Science	0	0	0.0%	0	0.0%	100,000,000	0	0.0%	0	0.0%
CSC249ADSE16	Mark S. Gordon	Chemistry	0	0	0.0%	0	0.0%	14,000,000	47,477	0.3%	47,477	100.0%
CSC249ADSE17	Charlie Catlett, Scot..	Computer Science	1,500,000	0	0.0%	0	0.0%	0	0	0.0%	0	0.0%
CSC249ADSE22	Zhenyu Huang	Energy Technologies	5,000,000	1,284,654	25.6%	123,795	9.6%	0	0	0.0%	0	0.0%
CSC250STDM10	Suren Byna	Computer Science	230,000	1,195,159	519.6%	196,408	16.4%	2,000,000	0	0.0%	0	0.0%
CSC250STDM11	Scott Klasky	Computer Science	230,000	412,785	179.4%	251,831	61.0%	700,000	15,556	2.2%	0	0.0%
CSC250STMS07	Barry Smith	Mathematics	230,000	38,762	16.8%	0	0.0%	0	0	0.0%	0	0.0%
CSC250STMS15	Wayne Joubert	Mathematics	230,000	0	0.0%	0	0.0%	700,000	0	0.0%	0	0.0%
CSC250STPM07	Jeff Daily	Computer Science	128,000	13,782	10.7%	0	0.0%	0	0	0.0%	0	0.0%
CSC250STPM11	Jack Dongarra, Earl..	Computer Science	0	0	0.0%	0	0.0%	700,000	472,622	67.5%	0	0.0%
CSC250STPM17	Scott Baden	Computer Science	128,000	311,074	243.0%	0	0.0%	32,000	2,442	7.6%	0	0.0%
CSC250STTO09	Jack Dongarra, Earl..	Computer Science	230,000	534	0.2%	0	0.0%	0	0	0.0%	0	0.0%
CSC250STTO11	John Mellor-Crumm ..	Computer Science	50,000	1,852	3.7%	0	0.0%	50,000	0	0.0%	0	0.0%
		Totals:	**129,736,000**	**38,191,781**	**29.4%**	**8,091,918**	**21.1%**	**379,182,000**	**139,400,911**	**36.7%**	**71,792,686**	**51.5%**

- Theta Capability: Jobs that used at least 20% of the machine
- Mira Capability: Jobs that used at least 16.7% of the machine

Figure 3.14 Exascale Computing Project Website Screen Shot

The ALCF helped to promote ECP training events to its users throughout the year in coordination with the ECP training lead. Tim Williams, then acting as ALCF's deputy director of science, also spoke to the training lead about possible future joint ECP-ALCF training events. ALCF Deputy Director of Science Kalyan Kumaran worked with the ECP Software Technology Director, Rajeev Thakur, to identify people at the E6 labs — the six national laboratories that are part of the ECP and have HPC facilities — who are working on ECP software technology projects and require access to the simulator at ALCF and helped provide those individuals with the access upon their request.

ECP Director Doug Kothe periodically holds teleconferences with the directors of the facilities, and ALCF Division Director Michael E. Papka regularly participates in those calls. Lastly, Katherine Riley participated in the capacity of reviewer (red team and actual review) for a couple of ECP reviews in the fourth quarter of CY 2017.

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 211 papers in CY 2017. 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 2017: 4.5 billion to INCITE, 1.7 billion to ALCC, and 950 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.

Section 4. Innovation

(a) Have innovations been implemented that have improved the facility's operations?

(b) Is the facility advancing research, either intramurally or through external collaborations, that will impact next generation high performance computing platforms, consistent with the DOE strategic goals? Is the facility tracking research results from the broader community, evaluating strategies to incorporate them into next generation high performance computing platforms?

ALCF Response

Listed below are the innovations and best practices carried out at ALCF during CY 2017. 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 Operational Innovation

4.1.1 PoLiMEr: An Energy Monitoring and Power Limiting Interface for HPC Applications

We have developed PoLiMEr, a user-space library for power monitoring and capping of HPC applications. PoLiMEr enables HPC users to navigate the power monitoring options on their system, to profile large-scale HPC applications, and to explore power-capping effects using intuitive and easy-to-use interfaces. PoLiMEr exhibits extremely low overhead and has successfully scaled to 1,024 nodes on Theta. PoLiMEr fully harnesses a system's power monitoring and control capabilities by utilizing all available power monitoring and control interfaces and providing an overview of all measurable components without requiring users to know what these power domains are. Simultaneously, through an automated process and thread detection, PoLiMEr can profile applications that use any combination of processes and threads. PoLiMEr's power monitoring and control functionalities enable users to monitor the power and energy consumption levels of their applications over time, profile specific portions of code, and set power caps at any time during the application's execution and on any combination and scale of nodes. PoLiMEr is being released as open-source and is already being used by other projects on Theta.

4.1.2 RAM Area Network

As reported last year, the RAN (RAM Access Network) project enables research into deep memory hierarchies and different cluster memory configurations. ALCF has a 5.5-TB pool of RAM available over the InfiniBand (IB) network, with 14 FDR (Fourteen Data Rate) IB 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 in future systems. The latencies are typically within 1 microsecond of the raw IB latencies and have performed extremely consistently in testing so far. This low amount of jitter should enable consistent performance, and for streaming, 50 percent of the local memory bandwidth capacity has been achieved 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. A number of interfaces to the memory pool are available, including malloc, mmap, memkind, block, and QEMU/KVM/RHEV, as well as a pmem-io persistent memory interface, a simple MPI-I/O interface, and a direct C application programming interface (API) that enables maximum performance by bypassing the kernel.

New this year, in terms of infrastructure, is provisioning, which is the ability to ask for more RAM than is on one memory appliance, and have it transparently appear as a single memory space. This functionality, provided by Kove, is currently in beta testing on our system. We have also investigated a number of application spaces. Our initial work with the GAMESS electron structure code has shown us numerous performance issues, which has led to significant improvements. (A poster on this work was presented at SC17.) Some initial work has been done with QMCPack. In this case, the need for some required code refactoring has been identified to hide memory access latency so as to achieve optimal performance. These two examples used the Kove C APIs, which is the highest-performing interface but requires changes to the application code.

Our most exciting application area, however, has been machine learning (ML). Three students worked with ALCF as part of a summer internship to use machine learning to address genomics problems via standard Python ML libraries. Here the transparent malloc interface was used, which required no changes to the code, only an LD_PRELOAD of the Kove malloc driver. A variety of algorithms were employed across these problems, including support vector machines, random forest, gradient boosting, long short-term memory, region-based convolutional neural networks, variational auto-encoders, and deep residual networks. Consistently, across all of these algorithms and science goals, results showed that as long as 5 percent to 20 percent of the RAM was local to the machine, with the rest in the remote RAN pool, the time to solution was essentially the same as when all of the RAM was local. These tests represent very early data, and more testing is needed, but the result indicates that ML/DL (deep learning) algorithms have the potential to greatly benefit from this technology. We are in the process of writing up a proposal for a graduate thesis project to further investigate/confirm this finding.

4.1.3 Operational Improvements

4.1.3.1 Audience – A File System Performance Monitoring Tool

Audience is a web application developed by ALCF staff member Zach Nault designed to monitor file system traffic and performance. It presents reads/writes per second and file opens/closes per second in graphical form for live and historical GPFS cluster performance. Audience is also able to provide performance at a job level both historically and live.

Audience is organized into three sections. The first is the Compute Resource view that shows top-level performance for all GPFS clusters connected to a given resource. The second is a Job view that presents a list of all active jobs running on Mira, Cetus, or Vesta, and specifically that job’s live file system performance. The last is a Historical view that provides the ability to query for performance data on a given cluster and/or over a given time frame, set of nodes, set of file systems, or performance metrics. We could also view a job's file system performance over the lifetime of the job, given the job ID.

The ALCF team has already used Audience to troubleshoot issues encountered with a specific project team's jobs that caused congestion within the ALCF storage network. The tool was able to provide a visual representation of the number of files being opened and closed during the life cycle of specific jobs, thereby identifying a bottleneck that could be avoided by changing some file system settings. This tool is currently in beta testing mode within the ALCF.

4.1.3.2 On-Demand Rebooting on Theta

Theta's KNL processors have some unique memory capabilities that users can take advantage of to improve their code performance. The node has two kinds of memory available, a small amount of high-bandwidth memory (HBM), along with a larger amount of available DDR (double-data-rate) RAM for use with their applications. Users may choose to use the HBM as a cache for their DDR RAM and may gain performance with no change to their code, or, if appropriate to their code and workload, they may choose to treat the HBM as a separate memory space and manually optimize their code to achieve even greater memory gains. In addition, users may change the non-uniform memory access (NUMA) mapping of their memory, providing better memory locality and better access times to KNL cores. Careful use of both of these features may provide a user with a very significant boost in performance. However, the optimal combination of NUMA and HBM modes can vary widely from code to code based on constraints from their solvers and science domains. In addition, changing these parameters on the KNL requires a costly node reboot operation.

We have enabled on-demand node rebooting with Cobalt on Theta, allowing users to request various NUMA and HBM mode configurations on a per-job basis, and Cobalt can now make decisions in job placement based on the requested memory characteristics of a job and current memory status on Theta's nodes. This mode has allowed us to run without a "default" node configuration on the system and allows us to mitigate the long node reboot times on Theta. The reboot mechanism also allows us to carefully track both (1) the amount of time that Theta spends rebooting nodes due to user requests and (2) improvements as Cray provides revisions to its software stack. This mechanism also allows users to choose and use the most efficient mode for their code(s) without any intervention on the part of the site system administrators. This feature was made available in the Early Science pre-production phase of bringing up Theta and was ready in advance of ALCC production, allowing users to use these modes as they saw fit for their projects from day one.

4.1.3.3 ALCF Publication Discovery Program

The ALCF Publication Discovery Program (ALCF PDP) is an automatic program for finding academic research articles published by ALCF users and referencing the Argonne contract number. The ALCF PDP is written in Python and uses web search APIs provided by Crossref, a Digital Object Identifier (DOI) registration agency of the International DOI Foundation. It uses information about users retrieved from the ALCF's Business Intelligence data warehouse as search criteria, along with the Argonne contract number for user facilities. This method casts a wider net than searching by using the ALCF only, as Argonne operates two other DOE/SC user facilities. Subsequent to this discovery phase, ALCF staff members inspect, tag, and verify each search result before adding it to the publications database on the ALCF website. The PDP program also uses the publication links from the search results to download the publications (.PDF files, etc.). All results from PDP are stored within the ALCF's data warehouse for easy access to all ALCF staff and to enable quick reporting.

4.2 Research Activities for Next Generation Systems

4.2.1 Joint Laboratory for System Evaluation

Argonne's Joint Laboratory for System Evaluation (JLSE) enables researchers to assess and improve next-generation computing platforms of interest to the DOE. Established by the ALCF and the MCS division, the JLSE centralizes Argonne's research activities aimed at evaluating future extreme-scale computing systems, technologies, and capabilities. JLSE users from the ALCF and MCS divisions leverage existing infrastructure and next-generation hardware and software to explore low-level experimental computer and computational science, including operating systems, messaging, compilers, benchmarking, power measurements, I/O, and new file systems. By providing access to leading-edge computing resources and fostering collaborative research, the JLSE enables researchers to address Argonne and DOE needs in a variety of areas, including:

- Improving science productivity on future hardware and software platforms.
- Providing an avenue for Argonne researchers to work collaboratively with HPC vendors on prototype technologies for petascale and beyond.
- Investigating alternative approaches to current and future system deployments.
- Maintaining a range of hardware and software environments for testing research ideas.
- Helping to drive standards on benchmarks, programming models, programming languages, memory technologies, etc.

The joint laboratory provides users with access to several diverse testbeds, including: An Intel Xeon Phi (Knights Landing) and Xeon (Skylake) cluster; an IBM Power System S822LC; an NVIDIA DGX-1 System (Volta GPUs); an Applied Micro X-C1 Server ARM Development Platform; an Atos Quantum Learning Machine; an IBM Elastic Storage Server GL6; and a Lustre testbed.

In 2017, the JLSE supported 374 users participating in 44 projects. These projects ranged from application portability to software development to tools and compiler development for an ALCF Early Science Project. The following summaries represent a sampling of current JLSE projects:

ALCF Data Science Program: Research teams from the ALCF Data Science Program use JLSE resources to explore and improve data science techniques, such as data mining, graph analytics, machine learning, and complex and interactive workflows.

ALCF Early Science: Research teams from the ALCF Early Science Program use JLSE resources to prepare and optimize applications for the next-generation supercomputers in advance of the systems being available. For example, researchers from the Theta ESP used JLSE's Intel Xeon Phi cluster to work on single node optimization, focusing on memory modes and vectorization.

Argo: Argo is a new exascale operating system and runtime system designed to support extreme-scale scientific computation. Researchers from the Argo project used JLSE resources to prototype the GlobalOS distributed resource management and to evaluate the performance of NodeOS. They also used the laboratory's testbeds to develop and optimize a lightweight, low-level threading and task framework for OpenMP and other programming models (Cilk, Quark, Charm++).

Big Data: Researchers are using JLSE testbeds to study the layering of HPC programming models beneath big data programming models. Specifically, they are researching the development of a software environment with a Spark user interface (Java and Scala) that can run on a supercomputer, cluster, or cloud with a back end for executing data-intensive communication patterns.

CANDLE: Using the NVIDIA DGX-1 system and other JLSE computing resources, researchers are developing the CANcer Distributed Learning Environment (CANDLE), a computational framework designed to facilitate breakthroughs in the fight against cancer.

Deep Learning: Multiple projects are using JLSE systems to investigate the potential of deep learning. One research team is focused on understanding how deep learning can be used to improve lossy compression of scientific data from simulations and instruments. Another effort is exploring the performance of different machine learning frameworks that have implemented deep learning and neural networks on KNL systems.

LLVM: Researchers used the JLSE's IBM Power systems to advance low-level virtual machine (LLVM) compiler development. The instruction set architecture for these systems is the same as for the IBM Blue Gene/Q system, with the only difference being in vectorization. LLVM and CLANG builds were carried out on the Intel Xeon Phi systems for quality assurance purposes. Researchers can complete these builds in 10 minutes using JLSE resources (compared to hours on a laptop).

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

Quantum Computing: A research team is using the JLSE's Atos Quantum Learning Machine and other resources to develop an artificial neural network for spectral analysis called Spectranne. This tool will automate the analysis of vast amounts of data being produced by state-of-the art, chirped-pulse spectroscopy experiments.

TAU: Argonne's Tuning and Analysis Utilities (TAU) performance system was ported to various instruction set architectures and other JLSE systems, including AppliedMicro's X-Gene ARM64, Intel Xeon Phi clusters, and IBM Power systems. Researchers validated support for OpenMP profiling via the OpenMP tools interface and vectorization intensity on the JLSE systems.

The JLSE also presented 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.

In addition, the JLSE hosted training opportunities such as the "Mini-workshop on Intel FPGA Technology in High Performance Computing," held July 26, 2017. This workshop was led by Altera/Intel technical specialist Greg Nash and focused on the use and programmability of Intel field-programmable gate arrays (FPGAs) for HPC. The OpenCL interface and associated OpenCL compiler were also discussed.

4.2.2 HPC Standards, Benchmarks, and Technologies

ALCF researchers are actively involved in developing standards, benchmarks, and technologies that help drive continued improvements in supercomputing performance. Staff activities include contributing to the C++ Standards Committee (Hal Finkel); the OpenMP Architecture Review Board (Kalyan Kumaran); the OpenMP Language Committee (Ray Loy); the OpenFabrics Interfaces Working Group (Paul Coffman); the HPCXXL User Group (Loy); the Standard Performance Evaluation Corporation (Kumaran); and the ACM SIGHPC Systems Professionals chapter (William Scullin). ALCF staff members also serve on the boards of Open Scalable File Systems, Inc. (Kevin Harms), a not-for-profit dedicated to supporting the Lustre file system; and the LLVM Foundation (Finkel), a not-for-profit focused on supporting the LLVM project and advancing the field of compilers and tools.

4.3 Best Practices

4.3.1 Best Practice: Library Tracking

As mentioned in last year's report, understanding the usage of ALCF resources enables us to focus our efforts in ways that best support our users. For instance, knowing which software libraries are actively in use by end user applications not only helps to prioritize library development and maintenance efforts, but also proves useful when planning future systems. For Mira, this information is collected using the ALCF-developed TrackLib tool and stored in the ALCF data warehouse. Since TrackLib was developed specifically for the Blue Gene/Q, significant effort would be required to port it to the Cray environment. Instead, to obtain this information for our latest system, Theta, we are using the XALT library tracking system, which has been in use on Cray systems for many years, and we have customized it to meet the specific needs of the ALCF. XALT provides the ALCF not only with the libraries an application uses and the frequency with which that application is used, but also with information that is useful when analyzing job failures. XALT went into production in June and is publishing to the data warehouse. A single interface for interrogating data from both Mira and Theta is under development and is anticipated to be available in 2018.

4.3.2 Best Practice: Software Refresh of Discretionary Allocations Management Website

In keeping with industry best practices, ALCF updated the website used to manage the DD allocation requests from users. The new website provides decision makers with the ability to filter and search for requests and provides other stakeholders with a mechanism to check on the status of previous requests. It also includes features such as annotating requests and saving partially filled-out responses, which allows for better remote collaboration among the Allocations Committee members.

4.3.3 Best Practice: Studying Performance Impact on ALCF Applications

Maintaining the ratio of network injection bandwidth (the rate at which a node injects data into the network fabric) to computational FLOPs seen in today's leadership-class systems is challenging in terms of cost, power, and physical size, especially as we look at the design of exascale computing architectures. ALCF performance engineers evaluated the overall performance impact of decreasing network injection bandwidth relative to increasing computational FLOPs. Using MPI data collected by the ALCF performance tool AutoPerf, ALCF staff projected the network injection bandwidth impact on exascale time to solution for the most heavily used scientific simulation codes at the ALCF.

4.3.4 Best Practice: Understanding and Reducing Run-to-Run Variation

ALCF performance engineering team members have been studying the performance variability on Theta as a factor of the different components of hardware. ALCF staff worked with Intel to understand the performance repeatability of these different components and how to improve the repeatability. The results of the study were published as a paper at SC17.

This page intentionally left blank.

Section 5. Risk Management

Is the Facility effectively managing operational risks?

ALCF Response

ALCF has clearly demonstrated successful risk management in the past year for both project and operation risks (Section 5.1). 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 **two** 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; and (2) Staff Recruitment Challenges, which is managed by ongoing recruiting and retasking of current staff as needed. The major risks tracked for the past year are listed in Section 5.2, with the risks that occurred and the mitigations for those risks described in greater detail in Table 5.1 and Section 5.3, respectively, along with one retired risk (Section 5.4, Table 5.2) and two new or recharacterized risks (Section 5.5, Table 5.3), as well as the major risks that will be tracked in CY 2018 (Section 5.6, Table 5.4).

Of primary interest here is a description of the most significant operational risks and the Risk Management Plan's effect on the Facility's day-to-day operations.

The Facility should provide an overview of:

- *The 3-5 most important operational risks that were tracked for the review year;*
- *Any significant new or re-characterized operational risks since the last review; and the operational risks that were retired during the current year;*
- *The major risks that will be tracked in the next year, with mitigations as appropriate.*
- *Any risk events that occurred, how the Risk Management Plan was implemented, and the effectiveness of mitigations.*

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 and 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 through the ALCF online risk reporting form.

5.1.1 Continuation of the ALCF-3 Project

The ALCF-3 project — procuring and deploying the next ALCF supercomputer — continued in CY 2017. 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.

5.1.2 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 strategies for the risk and considers whether any of the strategies need updating for any reason, including due to changes in the technology landscape.
  - Works with the risk owner to modify the risk statement should any risk information indicate a need for changes to risk mitigation strategies, risk triggers, or risk scope.

5.1.3 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 first half of CY 2017, operating both Mira and Theta during the second half of CY 2017, 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 2017.

Five major operations risks were tracked for CY 2017, **two** with a risk rating of High and three with a risk rating of Moderate. Of these, three were encountered and managed. The **five** 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 2017

ID	Title	Encountered	Rating	Notes
25	Staff recruitment challenges	Yes	High	ALCF added 6 new hires overall this year, plus 2 internal reclassifications. 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 **3** staff members during CY 2017. 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	Moderate	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.

5.3 Risks Encountered in the Review Year and Their Mitigations

The six risks encountered during CY 2017 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, and 3 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 the highest-scored risks encountered. The facility was required to operate with moderate austerity measures during the early part of the year due to the Continuing Resolution. 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 2017; however, the funding uncertainty was large enough to also require some budget reprioritizing of purchases and new hires.

Management

In conversation with DOE-ASCR and through budget deep dive activity, 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 2018 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. Conduct 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 2017 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: Inability 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, **three** ALCF staff members left, one of whom moved to a position elsewhere at Argonne. **Four** new full-time regular staff and four long-term temporary staff (two external hires, two reclassifications from postdocs) were added during CY 2017, for a net **gain of five** ALCF staff members for the year. Thus, while ALCF has continued to make good progress on adding new hires, staff retention remains a concern, as does hiring enough qualified staff to meet the staffing needs that now encompass operating both Mira and Theta.

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 demand continues to be high for the skills of ALCF staff members. Thus, both staff recruitment and staff retention will remain a focus for ALCF.

5.3.3 Problems with Water Cooling

1065: Problems with Water Cooling	
Risk Owner	Mark Fahey
Probability	Very Low
Impact	Cost: Very Low; Technical Scope: Very Low
Risk Rating	Very Low
Primary Management Strategies	Work with vendor to ensure that equipment is appropriately calibrated and monitored. Meet with Theory and Computing Sciences building management to assure that the management staff understands the system operations. Maintain a second process loop for redundancy. Monitor and issue notifications of possible issues.
Triggers	Regression testing indicates issue. User complaints.

Description

On February 6, 2017, water temperatures in the cooling loop increased to about 5 degrees Fahrenheit higher than normal. This temperature increase did not cause any problems for Mira.

Evaluation

Differential pressure across the Victaulic strainers was too large because of particulate clogging. For about two weeks before the incident, the chilled water plant had been compensating for the lower flow by lowering the chilled water temperature 2 degrees Fahrenheit, which resulted in the building cooling loop water also being lowered 2 degrees Fahrenheit. This temperature reduction brought the cooling loop temperature close enough to the target temperature that the clogging of the strainers was not noticed initially.

Management

The strainers were cleaned, which restored the pressure and flow to normal and thus restored the water temperature to normal.

5.3.4 Delays in Completing the Signing of Some INCITE/ALCC User Agreements

1012: Delays in Completing the Signing of Some INCITE/ALCC User Agreements	
Risk Owner	Richard Coffey
Probability	Very Low
Impact	Cost: Moderate; Technical Scope: Very Low
Risk Rating	Very Low
Primary Management Strategies	Put in place a series of reminders and outreach activities to members and their PIs. Status user agreements with the DOE Program Office. Track which projects have agreements signed.
Triggers	Incomplete set of signed agreements shortly before the start date of the INCITE or ALCC program year. Account processing for INCITE or ALCC project members that takes much longer than expected. Difficulties reaching PIs.

Description

One of the 2016–2017 ALCC projects still had not signed its MUA by early 2017. This delay meant that the project had accomplished nothing with the allocation from July 2016 through early 2017.

Evaluation

There was lack of communication between the PI and the legal department of the multinational organization at which the PI is employed. E-mail messages from both the PI and the legal department to ALCF gave ALCF the impression that the signing of the agreement was imminent, when in fact no progress was being made.

Management

In early 2017, the ALCF Director of User Experience and the Industry Outreach Manager brought the PI and the legal department together to discuss how to move forward and have the MUA signed. Because of all the time that had been lost, ALCF also contacted the DOE ALCC program manager and discussed the option of extending the allocation for an additional six months, to December 2017. The program manager approved this arrangement, and ALCF decided to take these hours from the DD allocation program.

The PI and the company legal department accepted this extension; the company then assigned a new person to work on the agreement, and the agreement was signed and in place within one month. During this time, ALCF worked to prepare the project so that the PI and his collaborators could immediately start using the allocation once the MUA was in place.

The lesson learned from this situation was that ALCF needs to be involved from the start in communications between a PI and the legal department of the PI's organization and shepherd along the process of approving and signing the MUA. The ALCF will continue to impress upon PIs the importance of the MUA.

5.3.5 A Failure of One Minor Component Can Lead to a Cascade of Failures That Becomes a Major Outage

1051: A Failure of One Minor Component Can Lead to a Cascade of Failures That Becomes a Major Outage	
Risk Owner	John Reddy
Probability	Very Low
Impact	Cost: Very Low; Technical Scope: Low
Risk Rating	Very Low
Primary Management Strategies	Periodically evaluate systems for potential areas of concern. Develop recovery plans for critical systems. Minimize single points of failure.
Triggers	Review uncovers new single point of failure. Failure of components in area of poor redundancy.

Description
On October 27, 2017, an uncorrectable memory check error caused a key management host to fail. This error prevented the start of any new jobs, blocking use of Theta until repaired.

Evaluation
An uncorrectable memory check error caused MOM3 (one of 3 Cray MOM server nodes) to crash. This crash interrupted the Cobalt job scheduler and caused any jobs using MOM3 to fail. Although jobs using MOM1 or MOM2 for job coordination could complete successfully, no new jobs on Theta could start until MOM3 was repaired.

Management
The failed memory stick in MOM3 causing the memory check error was replaced, and MOM3 was restored to service. New jobs were able to start, and Theta was returned to normal operation.

5.4 Retired Risks

One risk was retired during the past year. This risk is summarized in Table 5.2.

Table 5.2 Risks Retired during CY 2017

ID	Title	Rating	Management Strategies	Notes
1087	Data analysis servers do not have sufficient memory	Very Low	Consider adding Solid State Drives (SSDs), additional software memory, out-of-core techniques.	The new visualization cluster that was brought online in 2015 has increased memory capacity. The two user groups that had previously experienced insufficient memory have not experienced this issue with the new cluster. The decision was therefore made to retire the risk.

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. In total, **1** new risk was added and **1** risk was recharacterized in CY 2017. Table 5.3 describes these risks.

Table 5.3 Risks That Were New or Recharacterized in CY 2017

ID	Title	Rating	Management Strategies	Notes
1081	Storage area network instability or performance problems	Very Low	The ALCF maintains storage area network testbeds for configuration and software testing and for hardware diagnostics; the ALCF has vendor support contracts in place to cover all failure modes; the ALCF maintains a replacement parts cache to lessen the impact of failures.	Restructuring of original risk 1081, recharacterization of the risk.
1105	Catastrophic failure of scratch file systems	Moderate	The ALCF will educate users on which filesystems are backed up, and which are ephemeral. The ALCF will deploy a hierarchical storage manager (HSM) on the vulnerable file systems.	New risk

5.6 Projected Major Operating Risks for the Next Year

Table 5.4 lists the current top operating risks projected for CY 2018, along with the current risk rating and management strategies for each 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 2018

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	Moderate	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 Integrated Safety Management 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 and shared cloud-based storage system, along with risk forms and a risk register that are both formatted using Excel. Over the past year, **1** risk was retired, **1** new risk was added, and **1** risk was recharacterized. Beyond these activities, many tools are used to manage risks at the ALCF, particularly in the area of safety. The ALCF's effective risk management plan has contributed to the successful management of all significant risks encountered in the past year.

This page intentionally left blank.

Section 6. Site Office Safety Metrics

(a) Does the Facility exhibit a culture of continual improvement in Environment, Safety, and Health (ES&H) practices to benefit staff, users, the public, and the environment?

(b) Has the Facility implemented appropriate ES&H measures?

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. Risk-based management assessments are conducted on an annual basis in order to ensure continual ES&H improvement. In CY 2017, assessments were conducted to review training compliance and use of personal protective equipment (PPE) in the data center. In addition, the ALCF performed a full evaluation and management assessment of its electrical work procedures. Work control documents were updated to reflect changes in Argonne's laboratory procedures, and the ALCF's Qualified Electrical Worker (QEW) qualifications and training were reviewed. The ALCF reminds workers annually, as well as during initial training, of their power and responsibility to stop work if an unsafe situation is encountered. All required safety training was completed by all workers prior to work being performed.

The ALCF employs appropriate work planning and control principles using Argonne's Work Planning and Control (WPC) and Work Authorization and Execution (WAE) systems. A formal WPC document is used for routine tasks such as racking/un-racking equipment, cabling/optical work, and working under a data center raised floor. Similarly, 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 (mild chemical hazard owing to water treatment chemicals, materials handling, 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. Whether routine or complex, work is authorized to be performed using the WAE system for each worker to ensure all workers have the required training and are aware of potential hazards and controls, including the required PPE for each task authorized.

In CY 2017, six "Mode 1" electrical jobs and 43 "Mode 3" electrical jobs were performed with no incident. The 43 Mode 3 electrical jobs also presented a mild chemical hazard that was not encountered. The identification in CY 2017 of a new Mode 1 electrical hazard — replacing the AC to DC power rectifiers in the XC40 compute racks — resulted in the creation and approval of a new WPC document. Despite the components being designed to be hot swapped, the ALCF determined there was sufficient risk of work injury or equipment damage to warrant a formal, complex WPC procedure. The ALCF worked with the vendor to define the procedure for zero energy verification to ensure that the equipment was in an electrically safe working condition. The ALCF also worked with the XC40 vendor to design sound dampening cabinets to reduce the ambient noise level around the machine blowers from 97 dBA to below 85 dBA. Following the standards of the American Conference of Governmental Industrial Hygienists (ACGIH), which are stricter than Occupational Safety and Health Administration standards for noise management, Argonne uses noise levels above 85 dBA to determine required PPE and possible inclusion in a hearing conservation program. Achieving this reduction meant that hearing protection is no longer required — although still encouraged — when working in the data center.

Section 7. Cyber Security

(a) Does the Facility exhibit a culture of continual improvement in cyber security practices?

(b) Have innovations been implemented to improve the Facility's cyber security posture?

(c) Does the Facility have a valid cyber security plan and Authority to Operate?

ALCF Response

The ALCF hired a dedicated Cyber Security Program Representative (CSPR) to continually monitor the ALCF security posture and to update policies and procedures to help ensure that leadership-class science can be conducted in a secure environment. This CSPR is matrixed into the Argonne Cyber Security Program Office (CSPO) to help ensure alignment with Argonne's security posture. As new user services were deployed in CY 2017, the ALCF proactively worked with the CSPO to perform risk analysis and to implement any additional necessary controls. The ALCF also improved its cyber security posture by working closely with Argonne's Cyber Operations, Analysis, and Research (COAR) team to implement a multi-factor authentication solution for Argonne. Level 4 Homeland Security Presidential Directive 12 (HSPD-12) cards have been, or will soon be, obtained by all privileged users of ALCF. Authentication using DOE-issued HSPD-12 badges is now required for all privileged users to access ALCF resources. A review of the ALCF cyber security plan was also completed, and all privileged users successfully completed annual Argonne cyber security training.

There were no cyber security incidents in CY 2017.

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

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? Are there other metrics that would have helped you assess the facility's operational performance?

ALCF Response

The ALCF and the DOE have agreed to the 2018 metrics and targets as proposed in the March 2017 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 2019, the proposed metrics and targets for the current production resources remain the same as for 2018.

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 2019, the proposed metrics and targets for the current production resource, Mira, will remain the same as for 2018. Appropriate metrics and targets have been proposed for Theta, a resource that entered production in 2017.

The 2018 metrics are covered in Section 8.2, and the 2019 metrics are covered in Section 8.3.

8.2 ALCF 2017 OA Performance Metrics

The Operational Assessment (OA) performance metrics, 2017 targets and actuals, and agreed-upon 2018 targets are presented in Table 8.1. The new ALCF resource, Theta, went into production in July 2017 and provided 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 2018 Scheduled Availability target for Theta will be 90 percent. Theta will provide INCITE hours during 2018; therefore, there is a new capability metric for Theta in 2018.

Table 8.1 Performance Metrics: 2017 Targets, 2017 Actuals, and Agreed-Upon 2018 Targets

Area	Metric	2017 Targets	2017 Actuals	2018 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	4.6/5.0	3.5/5.0
	User Survey – **Response Rate**	25.0%	46.9%	25.0%
	% User Problems Addressed Within Three Working Days	80.0%	95.8%	80.0%
Business Results	Mira Overall Availability	90.0%	96.4%	90.0%
	Mira Scheduled Availability	90.0%	99.5%	90.0%
	% of INCITE core-hours from jobs run on 16.7% or more of Mira (131,072–786,432 cores)	40.0%	72.2%	40.0%
	% of INCITE core-hours from jobs run on 33.3% or more of Mira (262,144–786,432 cores)	10.0%	27.5%	10.0%
	Theta Overall Availability (July 2017–December 2017)	80.0%	90.0%	90.0%
	Theta Scheduled Availability (July 2017–December 2017)	90.0%	98.6%	90.0%
	% of INCITE core-hours from jobs run on 20% or more of Theta (>= 51,328 cores)	N/A	N/A	20%

8.3 ALCF Proposed 2019 OA Performance Metrics

The Operational Assessment performance metrics, agreed-upon 2018 targets, and 2019 proposed targets are shown in Table 8.2.

Table 8.2 Performance Metrics: Agreed-Upon 2018 Targets and Proposed 2019 Targets

Area	Metric	2018 Targets	Proposed 2019 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 2018 Targets and Proposed 2019 Targets (Cont.)

Area	Metric	2018 Targets	Proposed 2019 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	90.0%	90.0%
	Theta Scheduled Availability	90.0%	90.0%
	% of INCITE core-hours from jobs run on 20% or more of Theta (>= 51,328 cores)	20.0%	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 2018 metrics and targets are reasonable measures of facility performance that are consistent with metrics and targets used at other facilities. For 2019, the proposed metrics and targets will remain the same as for 2018 for the current production resources. Achieving the agreed-upon 2018 and the proposed 2019 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 – Calculations

A.1 Scheduled Availability Calculation Details

Scheduled availability is the percentage of time a designated level of resource is available to users, excluding **scheduled outage** time 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 the regularly scheduled maintenance is not needed, users will be informed of the cancellation of the 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 the return to scheduled production will be counted as an adjacent unscheduled outage. Typically, this designation would be assigned 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.

Formula:

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

Where

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

A.2 Overall Availability Calculation Details

Overall availability is the percentage of time a system is available to users. Outage time reflects both scheduled and unscheduled outages.

***Formula*:**

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

A.3 System Mean Time To Interrupt (MTTI) Calculation Details

MTTI (Mean Time to Interrupt) is defined as time, on average, to any outage of the full system, whether unscheduled or scheduled. It is also known as MTBI (Mean Time Between Interrupts).

Formula:

$$\mathbf{MTTI} = \frac{\textbf{time in period} - (\textbf{duration of scheduled outages} + \textbf{duration of unscheduled outages})}{\textbf{number of scheduled outages} + \textbf{number of unscheduled outages} + \mathbf{1}}$$

A.4 System Mean Time To Failure (MTTF) Calculation Details

MTTF (Mean Time to Failure) is defined as the time, on average, to an unscheduled outage of the full system.

Formula:

$$\mathbf{MTTF} = \frac{\textbf{time in period} - \textbf{duration of unscheduled outages}}{\textbf{number of unscheduled outages} + \mathbf{1}}$$

A.5 Utilization Calculation Details

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. Jobs that ran during an outage are excluded.

Formula:

$$\mathbf{Utilization} = \left(\frac{\textbf{\textit{corehours used in period}}}{\textbf{\textit{corehours available in period}}}\right) * \mathbf{100}$$

A.6 Capability

Capability is an attribute assigned to user jobs that meet the capability definition for a machine. **High Capability** is an attribute assigned to user jobs that meet the high capability definition for a machine.

Tables A.1 and A.2 show the capability definitions for each reportable machine.

Table A.1 Capability Definitions for Mira

Mira				
Capability	High Capability	Range	Minimum Nodes/Cores	Maximum Nodes/Cores
No	No	0% <= x < 16.7%	1 / 16	8,191 / 131,056
Yes	No	16.7% <= x < 33.3%	8,192 / 131,072	16,383 / 262,128
Yes	Yes	33.3% <= x <= 100%	16,384 / 262,144	49,152 / 786,432

Table A.2 Capability Definitions for Theta

Theta				
Capability	**High Capability**	**Range**	**Minimum Nodes/Cores**	**Maximum Nodes/Cores**
No	No	0% <= x < 20.0%	1 / 64	647 / 41,408
Yes	No	20.0% <= x < 60.0%	648 / 41,472	1,943 / 124,352
Yes	Yes	60.0% <= x	1,944 / 124,416	See: A.7 Theta Nodes

Capability also refers to a calculation. The capability calculation is the percent of core hours of jobs with the capability attribute versus the total core hours of all jobs. The calculation can be applied to a class of jobs. For example: INCITE capability is the percent of core hours of INCITE jobs with the capability attribute versus the total core hours of all INCITE jobs for a time period.

***Formula*:**

$$\textbf{OVERALL CAPABILITY} = \left(\frac{\textbf{\textit{Capability Core Hours Consumed}}}{\textbf{Total Core Hours Consumed}}\right) * \textbf{100}$$

$$\textbf{HIGH CAPABILITY} = \left(\frac{\textbf{\textit{High Capability Core Hours Consumed}}}{\textbf{Total Core Hours Consumed}}\right) * \textbf{100}$$

A.7 Theta Nodes

The number of reportable nodes on Theta was fewer than the total nodes. The total node count for Theta changed during 2017, as shown in Table A.3.

Table A.3 Total and Reportable Nodes for Theta

Theta		
Data Range	**Total Nodes**	**Reportable Nodes**
07/01/2017 – 12/12/2017	3,624	3,240
12/13/2017 – 12/31/2017	4,392	3,240
01/01/2018	4,392	4,008

The reportable node count is used in the following calculations:

- Scheduled Availability: Affects the scheduled outage and unscheduled outage calculations when the node count in the outage was fewer than the total number of nodes.
- Overall Availability: Affects the scheduled outage and unscheduled outage calculations when the node count in the outage was fewer than the total number of nodes.

- Utilization: The calculation capped the daily utilization at 100 percent of reportable nodes. The number of core hours for each day was calculated as the minimum of the core hours used and the core hours possible.
- Overall Capability: 20 percent of the reportable nodes.
- High Capability: 60 percent of the reportable nodes.

Appendix B – ALCF Director's Discretionary Projects: Mira

January 1, 2017–December 31, 2017
Mira

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
3D_MOC_Reactor_Sim	John Tramm	Massachusetts Institute of Technology (MIT)	Domain Decomposed 3D Method of Characteristics Reactor Simulation Study	Nuclear Energy	1,000,000
Abeta42	Georges Belfort	Rensselaer Polytechnic Institute (RPI)	Investigating a Drug Target of the Amyloid Beta (Aβ) Peptide for Alzheimer's Disease with Molecular Simulations	Biological Sciences	4,000,000
AEM_Hydroxide_NYU	Mark E. Tuckerman	New York University	First-Principles Discovery of Design Rules for Anion Exchange Membranes with High Hydroxide Conductivity	Chemistry	5,000,000
ALCF_Getting_Started	Ray Loy	Argonne National Laboratory	ALCF Getting Started	Training	200,000
AlgLA	Oded Schwartz	The Hebrew University	Algorithmic Linear Algebra	Computer Science	250,000
Allinea	Ray Loy	Argonne National Laboratory	Improved Debugging Memory Usage for Blue Gene/Q	Internal	1,000,000
Alloy-Corrosion	Hendrik Heinz	University of Colorado-Boulder	Understanding and Designing Corrosion-Resistant Alloys	Materials Science	2,000,000
Angora_scaling_study	Allen Taflove	Northwestern University	Angora Scaling Study	Biological Sciences	9,980,145
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
ATPESC17_Instructors	Marta Garcia Martinez	Argonne National Laboratory	Argonne Training Program on Extreme Scale Computing for ALL Instructors	Training	100,000
ATPESC2017	Marta Garcia Martinez	Argonne National Laboratory	Argonne Training Program on Extreme-Scale Computing 2017	Training	10,000,000
aurora_app	William Scullin, Kevin Harms	Argonne National Laboratory	Aurora Application Enablement	Computer Science	8,200,000
AutomatedBench	Kevin Harms	Argonne National Laboratory	Automated Benchmarking	Computer Science	2,000,000
BIP	Pratul Agarwal	Oak Ridge National Laboratory (ORNL)	Biophysical Characterization of Protein Landscape	Biological Sciences	20,000,000
bloodflow_dd	Jifu Tan	Northern Illinois University (NIU)	Microfluidic Design and Optimization for Cell Separation	Engineering	1,000,000
BrainImagingADSP	Doga Gursoy	Argonne National Laboratory	Large-Scale Computing and Visualization on the Connectomes of the Brain	Biological Sciences	10,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
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	12,000,000
Bubblecollapse	Eric Johnsen	University of Michigan	Dynamics and Heat Transfer in Bubble Collapse near Solid Surfaces	Engineering	4,500,000
Catalyst	Katherine Riley	Argonne National Laboratory	Catalyst	Internal	20,000,000
CEED_ECPAD	Misun Min	Argonne National Laboratory	CEED	Mathematics	2,000,000
CG_IMP_Integration	Thomas F. Miller III	California Institute of Technology (Caltech)	A New Paradigm for the Rational Expression of Integral Membrane Proteins	Biological Sciences	1,500,000
CharmRTS	Laxmikant V. Kale	University of Illinois at Urbana-Champaign	Charm++ and Its Applications	Computer Science	2,500,000
Cluster_catalysis	Anastassia N. Alexandrova	University of California-Los Angeles	Ensemble Representation for the Realistic Modeling of Cluster Catalysts at Heterogeneous Interfaces	Chemistry	2,500,000
CM1-LES	Gokhan Sever	Argonne National Laboratory	Large-Eddy Simulations of Airflow Dynamics and Physics around the Island of Graciosa	Earth Science	2,000,000
CMSHPCProd	Harvey Newman	California Institute of Technology (Caltech)	CMS Production on HPC	Physics	400,000
CMT	Scott Parker	Argonne National Laboratory	Compressible Multiphase Turbulence	Engineering	5,000,000
CNTmetallization	Iman Salehinia, Mike Papka	Northern Illinois University (NIU)	Metallization of CNT for Thermal and Structural Applications	Materials Science	1,500,000
ComplexTurbulence	Krishnan Mahesh	University of Minnesota	High-Fidelity Simulations of Complex Turbulent Flows	Engineering	7,000,000
Comp_Perf_Workshop	Ray Loy	Argonne National Laboratory	Comp_Perf_Workshop	Training	7,500,000
CORALDev	Scott Parker	Argonne National Laboratory	CORAL Development and Testing	Internal	3,000,000
CORALtestApps	James Osborn	Argonne National Laboratory	Preparing Test Applications for CORAL Machines	Internal	100,000
CosmicLaser_DD	Petros Tzeferacos	The University of Chicago (UChicago)	Simulations of Laser Experiments to Study the Origin of Cosmic Magnetic Fields	Physics	5,000,000
Cray	Ti Leggett, Mark Fahey, Susan Coghlan	Cray Inc.	Cray Installation	Internal	1,000,000
critical_perf	Scott Parker, Ray Loy	Argonne National Laboratory	Critical Debugging Project	Internal	40,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CrystalsADSP	Alexandre Tkatchenko, Alvaro Vazquez-Mayagoitia	University of Luxembourg	Constructing and Navigating Polymorphic Landscapes of Molecular Crystals	Materials Science	35,000,000
CWT	Justin M. Wozniak	Argonne National Laboratory	Cancer Workflow Toolkit	Biological Sciences	1,000,000
datascience	Venkatram Vishwanath	Argonne National Laboratory	ALCF Data Science and Workflows Allocation	Internal	2,000,000
DD-MacroDFT	Kaushik Bhattacharya	California Institute of Technology (Caltech)	Large-Scale *Ab-Initio* Simulations of Crystal Defects	Materials Science	4,000,000
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	100,000
dislocation-pattern	Tomas Oppelstrup	Lawrence Livermore National Laboratory (LLNL)	Strain Hardening and the Nature of Dislocation Patterns	Materials Science	2,000,000
DNS-ABF	Antonino Ferrante	University of Washington	Direct Numerical Simulation of an Aft-Body Flow	Engineering	1,000,000
DNSTFlow_PostProc	Jonathan Poggie	Purdue University	Direct Numerical Simulation of Compressible, Turbulent Flow	Engineering	5,000,000
DSEM_SUPERSONIC	Farzad Mashayek	University of Illinois at Chicago	Simulation of Supersonic Combustion	Engineering	6,000,000
duanl	Lian Duan	Missouri University of Science and Technology	Numerical Simulation of Acoustic Radiation from High-Speed Turbulent Boundary Layers	Engineering	3,000,000
EarlyPerf_aurora	Scott Parker	Argonne National Laboratory	Porting for Performance: The Software & Programming Environment Project to Enable High Performance Computational Science	Computer Science	20,000,000
EKaos	Ali Mani	Stanford University	DNS of Chaotic Electroconvection	Engineering	1,000,000
el-eff	Eva Zarkadoula	Oak Ridge National Laboratory (ORNL)	2T-MD Model Simulations of High Energy Ion Irradiation	Materials Science	2,200,000
em-brain	Mike Papka	Argonne National Laboratory	ANL-BrainProject	Biological Sciences	10,000
Enzo-PCello	Michael Norman	UC San Diego	Toward Exascale Hydrodynamic Cosmology	Physics	10,000,000
ExaHDF5	Venkatram Vishwanath	Argonne National Laboratory	ExaHDF5: Advancing HDF5 HPC I/O to Enable Scientific Discovery	Computer Science	2,000,000
EXCEL	Francisco Doblas-Reyes	Barcelona Supercomputing Center	EXtreme Climate Event Attribution Using Dynamical SeasonaL Predictions	Earth Science	630,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
ExM	Justin M. Wozniak	Argonne National Laboratory	Extreme Many-Task Computing with Swift	Computer Science	1,200,000
Extreme_Scale_TS	William M. Tang	Princeton Plasma Physics Laboratory (PPPL)	Extreme Scale Turbulence Simulations	Physics	5,000,000
FEP_REMD	Atsutoshi Okabe	Takeda Pharmaceutical Company Limited	Developing the Cutting-Edge Technology for Predicting Highly Accurate Binding Affinity between Protein and Ligand	Computer Science	2,000,000
Fermilab_Acc_PostPro	Eric G. Stern	Fermi National Accelerator Laboratory (Fermilab)	Fermilab Recycler Slip Stacking Simulations	Physics	10,000,000
FMO_EFMO_ECP	Federico Zahariev	Iowa State University / Ames Laboratory (DOE)	Fragment Molecular Orbital (FMO) and Effective Fragment Molecular Orbital (EFMO) Methods for Exascale Computing	Chemistry	6,000,000
fsi	Fande Kong	Idaho National Laboratory (INL)	Development of Nonlinear Eigenvalue Solvers in the Multiphysics Object Oriented Simulation Environment (MOOSE)	Mathematics	10,000,000
FTQCD	Xiao-Yong Jin	Argonne National Laboratory	Finite Temperature Lattice QCD Startup	Physics	2,000,000
fuel_rod_assemblies	Elias Balaras	The George Washington University	High-Fidelity Computations of Fuel Rod Assemblies under Seismic Loads	Engineering	5,000,000
Full_core_3D_MOC	Kord Smith	Massachusetts Institute of Technology (MIT)	Full Core PWR Simulation Using 3D Method of Characteristics	Nuclear Energy	23,889,917
GalaxiesOnFire	Philip F. Hopkins	California Institute of Technology (Caltech)	Galaxies on FIRE: Shedding Light on Dark Matter	Physics	2,000,000
Gas_Turbine_Cooling1	Sumanta Acharya	Illinois Institute of Technology (IIT)	Efficient Cooling Technologies for Advanced Gas Turbines and Power Plants	Engineering	1,000,000
GDB_Edge	Barrett N. Rogers	Dartmouth College	Global Two-Fluid Simulations of Tokamak Edge Region	Fusion Energy	5,000,000
GE_DualFuel	Sibendu Som	General Electric Company (GE)/Global Research	Computational Study of Cycle-to-Cycle Variation Dual-Fuel Engines	Energy Technologies	5,000,000
GLOBALFOUNDRIES_ prep	Sefa Dag	GLOBALFOUNDRIES Inc.	GLOBALFOUNDRIES Preparation for ALCC	Materials Science	250,000
GrainBoundaries	Wissam Saidi	University of Pittsburgh	Structure and Properties of Grain Boundaries in Materials for Energy Applications	Materials Science	2,000,000
HadronicLight_DD	Thomas Blum	University of Connecticut	Hadronic Light-by-Light Scattering Contribution to the Muon Anomalous Magnetic Moment from Lattice QCD with Chiral Fermions	Physics	527,747
HAP_Crystals	Iman Salehinia, Mike Papka	Northern Illinois University (NIU)	HAP_Crystals	Materials Science	500,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
HEPCloudFacility	Burt Holzman	Fermi National Accelerator Laboratory (Fermilab)	Testing the CondorCE Interface with HEPCloud Facility	Physics	200,000
HighReyTurb_PostProc	Robert D. Moser	The University of Texas at Austin	Data Analysis of Turbulent Channel Flow at High Reynolds Number	Engineering	7,000,000
HiPPSTR4096	Andrew D. Bragg	Duke University	High-Resolution DNS of Inertial Particle Motion in Turbulence	Engineering	1,800,000
HPCTuning	Khaled Ibrahim	Lawrence Berkeley National Laboratory	HPC Applications Tuning	Computer Science	5,000,000
hpMusic	Z.J. Wang	University of Kansas	A Study of Scalability of hpMusic	Engineering	1,000,000
Hybrid-C-Modelling	Jinxun Liu	U.S. Geological Survey (USGS)/Western Geographic Science Center (USGS)	Simulating Global Terrestrial Carbon Sequestration and Carbon Transport to Aquatic Ecosystems - Pilot Study	Earth Science	4,000,000
Injfuel	Marco Arienti	Sandia National Laboratories, California	High-Fidelity Characterization of Fuel Injection	Engineering	100,000
ITHALES	Koen Hillewaert	Cenaero	Industrial Turbomachinery High Accuracy LES – iTHALES	Engineering	4,000,000
LatticeQCD_aesp	Paul Mackenzie	Fermi National Accelerator Laboratory (Fermilab)	Lattice Quantum Chromodynamics Calculations for Particle and Nuclear Physics	Physics	10,000,000
LESOpt_DD	Qiqi Wang	Massachusetts Institute of Technology (MIT)	Adjoint Based Optimization via Large Eddy Simulation of a Fundamental Turbine Stator-Rotor	Engineering	5,000,000
LiquidGaNa_PostProc	Janet Scheel, Joerg Schumacher	Occidental College	Convective Turbulence in Liquid Gallium and Sodium	Engineering	3,000,000
Low_Pt_content_catal	Binay Prasai	Central Michigan University	Modeling and Prediction of Nanocatalyst for Fuel Cells	Energy Technologies	2,000,000
LPI_shlgnt	Jun Li	University of California-San Diego	Hot Electron Scaling and Energy Coupling in Nonlinear Laser Plasma Interactions	Physics	5,000,000
LQCDdev	James Osborn	Argonne National Laboratory	Lattice QCD Development	Physics	1,000,000
LTC_Aramco	Sibendu Som	Argonne National Laboratory/Aramco Services Company	Investigation of a Low Octane Gasoline Fuel for a Heavy-Duty Diesel Engine in a Low-Temperature Combustion Regime	Engineering	15,901,589
LUCMDproject	Nikolai Smolin	Loyola University Chicago	Molecular Dynamics Simulation of SERCA Troponin Complex and TRIM5a	Biological Sciences	368,640
Maintenance	Mark Fahey	Argonne National Laboratory	LCF Operations System Maintenance	Internal	714,000,000
MFPM	Olivier Desjardins	Cornell University	Large-Scale Simulation of Multiphase Flow in Porous Media	Engineering	1,999,526

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
mpich-bgq	William Scullin	Argonne National Laboratory	MPICH Testing on Blue Gene/Q	Computer Science	300,000
MPICH_MCS	Ken Raffenetti	Argonne National Laboratory	MPICH – A High Performance and Widely Portable MPI Implementation	Computer Science	26,000,000
NAMD_aesp	Benoit Roux	The University of Chicago (UChicago)	Free Energy Landscapes of Membrane Transport Proteins	Biological Sciences	5,000,000
nci_doe_pilot1	Rick Stevens	Argonne National Laboratory	nci_doe_pilot1	Biological Sciences	100,000
Nek_VHTR	Masahiro Kawaji	City College of New York (CCNY)	Scalability and Validation of Nek5000 for VHTR Challenge Problem of Pipe Flow Relaminarization	Engineering	2,000,000
Operations	Mark Fahey	Argonne National Laboratory	Systems Administration Tasks	Internal	42,000,000
OSCon	Andreas Glatz	Argonne National Laboratory	Optimizing Superconductor Performance through Large-Scale Simulation	Materials Science	8,000
Particle_Flow	Brian Helenbrook	Clarkson University	Examination of Particle-Wall Collision Models in Turbulent Particle Laden Flows	Engineering	1,812,355
PBSim	Lambert Fick	Texas A&M University (TAMU)	DNS of Incompressible Flows in Packed Pebble Beds	Engineering	2,000,000
Performance	Scott Parker, Ray Loy	Argonne National Laboratory	Performance	Internal	40,000,000
PHASTA_aesp	Kenneth Jansen	University of Colorado-Boulder	Extreme Scale Unstructured Adaptive CFD: From Multiphase Flow to Aerodynamic Flow Control	Engineering	1,000,000
PHASTA_NCSU	Igor A. Bolotnov	North Carolina State University (NCSU)	Multiphase Simulations of Nuclear Reactor Thermal Hydraulics	Engineering	12,000,000
PPBinding_PostProc	Benoit Roux	The University of Chicago (UChicago)	Finding Minimum Separation Pathway of Protein-Protein Complex	Biological Sciences	40,000,000
Practical_Strassen	Robert van de Geijn	The University of Texas at Austin	Make Strassens Algorithm Practical	Computer Science	50,000
PROTEUS_2016	Emily Shemon	Argonne National Laboratory	NEAMS Neutronics Verification and Validation Simulations	Nuclear Energy	7,000,000
Py8HPC	Stephen Mrenna	Fermi National Accelerator Laboratory (Fermilab)	Py8HPC	Physics	5,100,000
QETSc	Murat Keçeli	Argonne National Laboratory	Spectrum Slicing Eigensolver for *Ab Initio* Simulations	Chemistry	2,000,000
QMC-AE	Brenda M. Rubenstein	Brown University	Accurate Calculations of the Binding Energies of Dipole-Bound Anions	Chemistry	7,500,000
QMC-Alloying	Jeffrey C. Grossman	Massachusetts Institute of Technology (MIT)	Alloying Li Based High-Capacity Cathode Materials Using Quantum Monte Carlo	Materials Science	10,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
QMCwater	Dario Alfe	University College London (UCL)	Water at Surfaces from Highly Accurate Quantum Monte Carlo Calculations	Materials Science	2,000,000
QuantumDS	Alvaro Vazquez-Mayagoitia	Argonne National Laboratory	Quantum Mechanics and Data Science	Chemistry	4,000,000
Quinoa	Jozsef Bakosi	Los Alamos National Laboratory (LANL)	Asynchronous Navier-Stokes Solver on 3D Unstructured Grids for the Exascale Era	Engineering	4,000,000
radix-io	Philip Carns	Argonne National Laboratory	System Software to Enable Data-Intensive Science	Computer Science	2,000,000
rcduct	Hassan M. Nagib	Illinois Institute of Technology (IIT)	Influence on Duct Corner Geometry on Secondary Flow: Convergence from Duct to Pipe Flow	Engineering	3,000,000
ReaxCath	Christopher Knight	Argonne National Laboratory	Reactive Modeling of Battery Cathodes and Interfaces	Chemistry	4,000,000
rec_sironi	Lorenzo Sironi	Columbia University	Particle-In-Cell Simulations of Explosive Reconnection in Relativistic Magnetically-Dominated Plasmas	Physics	1,900,000
REEs_and_actinides	Deborah A. Penchoff	Institute for Nuclear Security	Accelerating Selective Binding of Rare Earth Elements and Actinides	Chemistry	3,936,314
REload_01	Francesco Grasso	Conservatoire National des Arts et Métiers (Le CNAM)	Aspect Ratio Effect in Closed and Open-Duct Flows	Engineering	4,000,000
RRMR_PostProc	Dmitri Uzdensky	University of Colorado	Kinetic Simulations of Relativistic Radiative Magnetic Reconnection	Physics	200,000
SciDAC_PILOT_Fission	Brian Wirth	The University of Tennessee Knoxville	Advancing Understanding of Fission Gas Behavior in Nuclear Fuel Through Leadership Class Computing	Nuclear Energy	3,500,000
SC_Architectures_DD	Tjerk Straatsma, Katerina B. Antypas, Timothy J. Williams	Oak Ridge National Laboratory (ORNL)	Portable Application Development for Next Generation Supercomputer Architectures	Computer Science	5,000,000
SDAV	Mike Papka, Joe Insley	Argonne National Laboratory	SciDAC Scalable Data Management Analysis and Visualization	Computer Science	1,000,000
SENSEI	Venkat Vishwanath	Argonne National Laboratory	Scalable Analysis Methods and *In Situ* Infrastructure for Extreme Scale Knowledge Discovery	Computer Science	2,000,000
SherpaProd	Taylor Childers	Argonne National Laboratory	Sherpa Parton Generator Production	Physics	5,000,000
SilentOwlFlight_DD	Anupam Sharma	Iowa State University (ISU)	Unraveling Silent Owl Flight to Develop Ultra-Quiet Energy Conversion Machines	Engineering	2,500,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
SkySurvey_PostProc	Katrin Heitmann	Argonne National Laboratory	Simulation for Large-Scale Cosmological Surveys	Physics	3,000,000
SolarWindowsADSP	Jacqueline M. Cole	University of Cambridge	Data-Driven Molecular Engineering of Solar-Powered Windows	Materials Science	197,000,000
STXQuantum	Vipin Sachdeva	Silicon Therapeutics	Comparison of Quantum Methods for Investigation of Protein-Ligand Binding Interactions	Chemistry	1,000,000
SU2_PadeOps_aesp	Sanjiva Lele	Stanford University	Benchmark Simulations of Shock-Variable Density Turbulence and Shock-Boundary Layer Interactions with Applications to Engineering Modeling	Engineering	300,000
Thermal_Transport	Lucas Lindsay	Oak Ridge National Laboratory (ORNL)	First Principles Thermal Transport for Basic Energy Sciences	Physics	500,000
Tools	Scott Parker	Argonne National Laboratory	ALCF Performance Tools	Internal	500,000
TopologyMapping	Zhiling Lan	Illinois Institute of Technology (IIT)	Topology Mapping of Irregular Applications	Physics	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	2,500,000
UH_polar	Julie McClean	Scripps Institution of Oceanography	Ocean and Sea Ice and Their Interactions around Greenland and Antarctica in Forced Fine-Resolution Global Simulations	Earth Science	1,940,000
umncafi	Joseph W. Nichols	University of Minnesota	Large Eddy Simulation for the Prediction and Control of Impinging Jet Noise	Engineering	5,000,000
UO2Behavior	Peter Littlewood, Laura Ratcliff	Argonne National Laboratory	Unveiling the Behavior of UO_2 under Extreme Physical Conditions	Materials Science	5,000,000
UofK_hypersonic_AMR	Christoph Brehm	University of Kentucky	Prediction of Boundary Layer Transition on Hypersonic Vehicles in Large Scale Wind Tunnels and Flight	Engineering	2,000,000
UrbanExP	Rajeev Jain	Argonne National Laboratory	Urban ECP	Engineering	1,400,000
UrbanLES	Robert Jacob	Argonne National Laboratory	LES Simulations of the Urban Boundary Layer	Earth Science	200,000
UTRC-Turbine	Chaitanya V. Halbe	United Technologies Research Center Inc.	High-Fidelity Simulation of Turbines in Engine-Relevant Conditions Toward Next-Generation Turbine Designs	Engineering	3,000,000
VarRhoFlow	Paul E. Dimotakis	California Institute of Technology (Caltech)	Variable-Density Fluid Dynamics	Engineering	10,000,000
VDVAT	Daniel Livescu	Los Alamos National Laboratory (LANL)	Variable-Density Turbulence under Variable Acceleration	Engineering	5,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
Vendor_Support	William E. Allcock	Argonne National Laboratory	Vendor Support	Internal	1,000,000
VirtualEye	Marco L. Bittencourt	University of Campinas (UNICAMP)	Computational Modeling of the Human Eye	Engineering	5,000,000
visualization	Michael E. Papka	Argonne National Laboratory	Visualization and Analysis Research and Development for ALCF	Internal	8,639,580
wall_turb_dd	Ramesh Balakrishnan	Argonne National Laboratory	Wall Resolved Simulations of Canonical Wall Bounded Flows	Engineering	5,000,000
WaterHammer	Hong Zhang	Argonne National Laboratory	Water Hammer Simulation	Mathematics	30,000
WOInterface	Baofu Qiao	Argonne National Laboratory	Interfacial Behavior of Alcohol at Water/Organic Biphasic System	Computer Science	5,000,000
				Total Mira DD	1,669,073,813

This page intentionally left blank.

Appendix C – ALCF Director's Discretionary Projects: Theta

July 1, 2017–December 31, 2017

Theta

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
21stcclimattrib_CESM	Samantha Stevenson	University of California-Santa Barbara	Attributing 21st Century Climate Change Impacts Using Targeted CESM Ensembles	Earth Science	3,000,000
Acceptance	Ti Leggett	Argonne National Laboratory	Acceptance Tests	Internal	10,000,000
AdvElecStructPr_tesp	Mark Gordon	Iowa State University (ISU)	Advanced Electronic Structure Methods for Heterogeneous Catalysis and Separation of Heavy Metals	Chemistry	10,000,000
AdvElecStruct_tesp	Mark Gordon	Iowa State University (ISU)	Advanced Electronic Structure Methods for Heterogeneous Catalysis and Separation of Heavy Metals	Chemistry	10,000,000
AD_Brain_Imaging	Jiook Cha	Columbia University	Computational Analysis of Brain Connectomes for Alzheimer's Disease	Biological Sciences	1,288,000
ALCF_Getting_Started	Ray Loy	Argonne National Laboratory	ALCF Getting Started	Training	10,000
Alzeta	Marc Day	Lawrence Berkeley National Laboratory	Improving Gas Reactor Design with Complex Non-standard Reaction Mechanisms in a Reactive Flow Model	Chemistry	1,700,000
AMS	S.C.C. Ting	Massachusetts Institute of Technology (MIT)	AMS Experiment MC Simulations	Physics	10,000,000
APSUpgradeLattice	Michael Borland	Argonne National Laboratory	Beam Dynamics for the APS Upgrade Lattice	Physics	4,000,000
AtlasADSP	Taylor Childers	Argonne National Laboratory	Advancing the Scalability of LHC Workflows to Enable Discoveries at the Energy Frontier	Physics	65,000,000
Atomization	Robert Saye	Lawrence Berkeley National Laboratory	HPC4Mfg: Modeling Paint Behavior During Rotary Bell Atomization	Engineering	1,000,000
ATPESC17_Instructors	Marta Garcia Martinez	Argonne National Laboratory	Argonne Training Program on Extreme Scale Computing for ALL Instructors	Training	100,000
ATPESC2017	Marta Garcia Martinez	Argonne National Laboratory	Argonne Training Program on Extreme-Scale Computing 2017	Training	2,000,000
AURORA	Ramesh Balakrishnan	Argonne National Laboratory	A Computational Study of a Proprietary Parallel, Arbitrary Cell Type Code Used by G.E. Global Research	Engineering	2,500,000
Banhammer_test	Haritha Siddabathuni Som	Argonne National Laboratory	Banhammer Test	Computer Science	10

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
Baryons_tesp	Katrin Heitmann	Argonne National Laboratory	Next-Generation Cosmology Simulations with HACC: Challenges from Baryons	Physics	100,000,000
BrainImagingADSP	Doga Gursoy	Argonne National Laboratory	Large-Scale Computing and Visualization on the Connectomes of the Brain	Biological Sciences	55,000,000
BrainModelingADSP	Fabien Delalondre	École Polytechnique Fédérale de Lausanne (EPFL)	Leveraging Non-Volatile Memory, Big Data, and Distributed Workflow Technology to Leap Forward Brain Modeling	Biological Sciences	36,000,000
BrainTissue_tesp	Fabien Delalondre, Tim Williams	École Polytechnique Fédérale de Lausanne (EPFL)	Large Scale Simulation of Brain Tissue	Biological Sciences	46,000,000
Candle_ECP	Rick Stevens	Argonne National Laboratory	Exascale Deep Learning and Simulation Enabled Precision Medicine for Cancer	Biological Sciences	1,000,000
Catalyst	Katherine Riley	Argonne National Laboratory	Catalyst	Internal	20,000,000
charles_theta	Frank Ham	Cascade Technologies, Inc.	Porting charLES onto KNL Architectures	Engineering	1,000,000
CharmRTS	Laxmikant V. Kale	University of Illinois at Urbana-Champaign	Charm++ and its applications	Computer Science	3,000,000
CM1-LES	Gokhan Sever	Argonne National Laboratory	Large-Eddy Simulations of Airflow Dynamics and Physics around the Island of Graciosa	Earth Science	250,000
Comp_Perf_Workshop	Ray Loy	Argonne National Laboratory	Comp_Perf_Workshop	Training	6,804,787
CONVERGE-SPP	Janardhan Kodavasal	Argonne National Laboratory	Scaling CONVERGE on KNL	Energy Technologies	100,000
CP2K_opt_KNL	Wei Jiang	Argonne National Laboratory	Porting and Optimization of Cp2k on Knights Landing Architecture	Computer Science	200,000
CP_SOD1	Shahar Keinan	Cloud Pharmaceuticals	Design of Novel Small Molecules to Inhibit SOD1 Aggregation as a Therapy for ALS	Biological Sciences	3,000,000
Cray	Ti Leggett, Mark Fahey, Susan Coghlan	Cray Inc.	Cray Installation	Internal	10,000,000
CrystalsADSP	Alexandre Tkatchenko, Alvaro Vazquez-Mayagoitia	University of Luxembourg	Constructing and Navigating Polymorphic Landscapes of Molecular Crystals	Materials Science	25,000,000
CSSI	Wei Jiang	Argonne National Laboratory	Scalable Reconstruction of X-ray Scattering Imaging for Nanomaterials	Physics	1,000,000
CWT	Justin M. Wozniak	Argonne National Laboratory	Cancer Workflow Toolkit	Biological Sciences	1,000,000
datascience	Venkatram Vishwanath	Argonne National Laboratory	ALCF Data Science and Workflows Allocation	Internal	10,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
DataSpaces	Manish Parashar	Rutgers University	DataSpaces Data Management Framework	Computer Science	100,000
DD-MacroDFT	Kaushik Bhattacharya	California Institute of Technology (Caltech)	Large Scale *Ab-Initio* Simulations of Crystal Defects	Materials Science	4,000,000
DE-SC0017033	Yueqin Huang	Cyentech Consulting LLC	Hardening of DOE UQ Codes for Ultra-Deep Proactive Geosteering in Oilfield Service	Earth Science	120,000
DeepLearningFromDFT	Ian Foster	The University of Chicago (UChicago)	Deep Learning from DFT	Materials Science	200,000
DUNE-LBNF	Nikolai Mokhov	Fermi National Accelerator Laboratory (Fermilab)	MARS Energy Deposition and Neutrino Flux Simulations	Physics	11,000,000
EarlyPerf_aurora	Scott Parker	Argonne National Laboratory	Porting for Performance: The Software & Programming Environment Project to Enable High Performance Computational Science	Computer Science	1,000,000
EarlyPerf_theta	Kalyan Kumaran	Argonne National Laboratory	Enabling Science on Theta: Tools, Libraries, Programming Models & Other System Software	Internal	10,000,000
ecp-testbed-01	Rajeev Thakur, Julia White, Doug Kothe, Bert Still, Pat McCormick, Susan Coghlan	Argonne National Laboratory	Exascale Computing Project (ECP) Testbed	Computer Science	37,420,231
EIC	Whitney Armstrong	Argonne National Laboratory	EIC Detector Simulations	Physics	100,000
Enzyme_Sim	Tao Li	Argonne National Laboratory	Stability and Activity Simulations of Biomass Enzyme under Complex Solvent Environment	Biological Sciences	5,000,000
ExaGrid	Michel Schanen	Argonne National Laboratory	Optimization and Dynamics Solvers for ExaGrid	Energy Technologies	50,000
EXASTEEL-THETA	Axel Klawonn, Oliver Rheinbach	TU Bergakademie Freiberg	Simulation of Dual-Phase Steel on Theta	Engineering	1,500,000
FFTBench	William Scullin	Argonne National Laboratory	FFT Benchmarking for Imaging Applications	Computer Science	750,000
FlameProp_tesp	Alexei Khokhlov	The University of Chicago (UChicago)	Direct Numerical Simulations of Flame Propagation in Hydrogen-Oxygen Mixtures in Closed Vessels	Chemistry	10,000,000
fpnmd2	Andre Schleife	University of Illinois at Urbana-Champaign	Qb@ll Intra-Node Scaling	Materials Science	5,000,000
FRC_Fusion_Modeling	Daniel Fulton	TAE Technologies, Inc.	Tri Alpha Energy FRC Whole Device Modeling – Phase I	Fusion Energy	5,000,000
FRNN	William Tang	Princeton University	Deep Learning for Fusion Energy Applications	Fusion Energy	4,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
Full_Core_TRRM	John Tramm	Massachusetts Institute of Technology (MIT)	High Fidelity Simulation of 3D Nuclear Reactor Cores Using The Random Ray Method	Nuclear Energy	3,000,000
FuncMats_tesp	Giulia Galli	The University of Chicago (UChicago)	First-Principles Simulations of Functional Materials for Energy Conversion	Materials Science	90,000,000
fvGFS	Jeffrey Durachta	Geophysical Fluid Dynamics Laboratory (GFDL)	Global Forecast System Using the Finite-Volume Cubed-Sphere Dynamical Core and GFS Physics	Earth Science	1,700,000
GA-bfg3d	Mark Kostuk	General Atomics	General Atomics Next Generation Skew-Symmetric Fluid Solver	Physics	150,000
GasJetsCylPrio_tesp	Christos Frouzakis	ETH Zurich	Flow, Mixing, and Combustion of Transient Turbulent Gaseous Jets in Confined Cylindrical Geometries	Chemistry	10,000,000
GasJetsCyl_tesp	Christos Frouzakis	ETH Zurich	Flow, Mixing, and Combustion of Transient Turbulent Gaseous Jets in Confined Cylindrical Geometries	Chemistry	10,000,000
GAtor	Noa Marom	Carnegie Mellon University	GAtor: A Cascade Genetic Algorithm for Crystal Structure Prediction	Materials Science	1,000,000
GFMCPrio_tesp	Steven C. Pieper	Argonne National Laboratory	Quantum Monte Carlo Calculations in Nuclear Theory	Physics	10,000,000
GFMC_tesp	Steven C. Pieper	Argonne National Laboratory	Quantum Monte Carlo Calculations in Nuclear Theory	Physics	10,000,000
GoAmazon	Usama Anber	Brookhaven National Laboratory	High Resolution Simulations of Deep Convection	Earth Science	50,000
HACC_aesp	Katrin Heitmann	Argonne National Laboratory	Extreme-Scale Cosmological Hydrodynamics	Physics	2,000,000
hifiturbfsi	Ivan Bermejo-Moreno	University of Southern California (USC)	High-Fidelity Simulation of Supersonic Turbulent Flow-Structure Interaction and Mixing	Engineering	2,000,000
HPC4Mfg_ACS	Deborah Bard	Lawrence Berkeley National Laboratory	Accelerating the Industrial Application of Energy-Efficient Chemical Separation	Materials Science	9,500,000
HybridPVPrio_tesp	Volker Blum, Noa Marom	Duke University	Electronic Structure Based Discovery of Hybrid Photovoltaic Materials on Next-Generation HPC Platforms	Materials Science	40,000,000
HybridPV_tesp	Volker Blum, Noa Marom	Duke University	Electronic Structure Based Discovery of Hybrid Photovoltaic Materials on Next-Generation HPC Platforms	Materials Science	10,000,000
HyperparADSP	Pierre Baldi	University of California at Irvine	Massive Hyperparameter Searches on Deep Neural Networks Using Leadership Systems	Computer Science	10,000,000
Intel	Kalyan Kumaran	Argonne National Laboratory	Intel Employees in Support of Theta	Internal	30,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
Job_Interference	Zhiling Lan	Illinois Institute of Technology (IIT)	Workload Interference Analysis on Theta	Computer Science	250,000
Landscapes	Michael Prentiss	Cambridge University	Exploring Protein Folding Energy Landscapes	Biological Sciences	20,000
LatticeQCD_aesp	Paul Mackenzie	Fermi National Accelerator Laboratory (Fermilab)	Lattice Quantum Chromodynamics Calculations for Particle and Nuclear Physics	Physics	5,000,000
LIGHTCONTROL	Sandra Biedron	University of New Mexico	Light Sources and Their Control Using AI Techniques	Physics	50,000
LPI_shlgnt	Jun Li	University of California-San Diego	Hot Electron Scaling and Energy Coupling in Nonlinear Laser Plasma Interactions	Physics	2,000,000
LSSTADSP_DESC	Katrin Heitmann	Argonne National Laboratory	Realistic Simulations of the LSST Survey at Scale - DESC	Physics	45,000,000
LSSTADSP_HACC	Katrin Heitmann	Argonne National Laboratory	Realistic Simulations of the LSST Survey at Scale - HACC	Physics	25,000,000
LUCMD2017	Nikolai Smolin	Loyola University Chicago	Molecular Dynamics Simulation of SERCA, SpeG, Troponin Complex, and TRIM5a	Biological Sciences	2,000,000
Maintenance	Mark Fahey	Argonne National Laboratory	LCF Operations System Maintenance	Internal	10,000,000
MembTrans_tesp	Benoit Roux, James C. Phillips	The University of Chicago (UChicago)	Free Energy Landscapes of Membrane Transport Proteins	Biological Sciences	60,000,000
MLResilience	Yanjing Li	The University of Chicago (UChicago)	Machine Learning Guided Cross-Layer Resilience	Computer Science	1,000,000
MMCM	Wei Jiang	Argonne National Laboratory	Multiscale Modeling for Catalysts and Materials	Biological Sciences	200,000
MOAB_App	Vijay Mahadevan	Argonne National Laboratory	MOAB Algorithmic Performance Portability	Mathematics	250,000
MueLu	Christopher Siefert	Sandia National Laboratories	Algebraic Multigrid in Trilinos	Mathematics	1,000,000
Multirefer-QMC	Anouar Benali	Argonne National Laboratory	Multireference Trial Wavefunction for Quantum Monte Carlo	Chemistry	2,000,000
MultiscaleADSP	Rathakrishnan Bhaskaran	GE Global Research	Enabling Multi-Scale Physics for Industrial Design Using Deep Learning Networks	Engineering	8,000,000
MultSimHem_dd	George Karniadakis	Brown University	Multiscale Simulations of Hematological Disorders	Biological Sciences	2,000,000
MuonMMPrio_tesp	Paul Mackenzie	Fermi National Accelerator Laboratory (Fermilab)	The Hadronic Contribution to the Anomalous Magnetic Moment of the Muon	Physics	10,000,000
MuonMM_tesp	Paul Mackenzie, Chulwoo Jung, Carleton DeTar	Fermi National Accelerator Laboratory (Fermilab)	The Hadronic Contribution to the Anomalous Magnetic Moment of the Muon	Physics	10,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
NAMD_aesp	Benoit Roux	The University of Chicago (UChicago)	Free Energy Landscapes of Membrane Transport Proteins	Biological Sciences	2,000,000
NAQMC_RMD_aesp	Aiichiro Nakano	University of Southern California (USC)	Metascalable Layered Materials Genome	Materials Science	10,000,000
NearSunTurb	Jean C. Perez	Florida Institute of Technology	Plasma Turbulence near the Sun	Physics	4,000,000
Nek_Tcc3	Saumil Patel	Argonne National Laboratory	Scalability of Grid-to-Grid Interpolation Algorithms for Internal Combustion Engine Simulations	Engineering	8,200,000
Nek_VHTR	Masahiro Kawaji	City College of New York (CCNY)	Scalability and Validation of Nek5000 for VHTR Challenge Problem of Pipe Flow Relaminarization	Engineering	500,000
networkbench	Sudheer Chunduri	Argonne National Laboratory	Network Benchmarking and Modeling	Computer Science	200,000
NWChemEx_aesp	Thomas Dunning, Alvaro Mayagoitia	Princeton Plasma Physics Laboratory (PPPL)	NWChemEx: Tackling Chemical, Materials & Biochemical Challenges in the Exascale Era	Chemistry	10,000,000
Operations	Mark Fahey	Argonne National Laboratory	Systems Administration Tasks	Internal	10,000,000
OSCon	Andreas Glatz	Argonne National Laboratory	Optimizing Superconductor Performance through Large-Scale Simulation	Materials Science	1,000,000
Performance	Scott Parker, Ray Loy	Argonne National Laboratory	Performance	Internal	10,000,000
PetaSupernovae	Sean Couch, Evan O'Connor	Michigan State University	Petascale Simulation of Magnetorotational Core-Collapse Supernovae	Physics	100,000
PetaSupernovae_DD	Sean M. Couch	Michigan State University	Petascale Simulation of Magnetorotational Core-Collapse Supernovae	Physics	100,000
PHASTA_aesp	Kenneth Jansen	University of Colorado-Boulder	Extreme Scale Unstructured Adaptive CFD: From Multiphase Flow to Aerodynamic Flow Control	Engineering	2,000,000
PICSSAR	Jean-Luc Vay	Lawrence Berkeley National Laboratory	Particle-In-Cell Scalable Spectral Relativistic	Physics	2,000,000
PKrylovPetsc	Oana Marin	Argonne National Laboratory	PETSc Pipeline Krylov	Computer Science	1,000,000
PRISMA	John N. Damoulakis	University of Southern California (USC)	Ptychography-based Rapid Imaging of Nano-Structures with Multilayer Assemblies	Materials Science	1,000,000
QETSc	Murat Keçeli	Argonne National Laboratory	Spectrum Slicing Eigensolver for *Ab Initio* Simulations	Chemistry	500,000
QE_KNL	Alvaro Vazquez-Mayagoitia	Argonne National Laboratory	Quantum Espresso in Theta	Materials Science	200,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
QMCCorrElec	Huihuo Zheng	Argonne National Laboratory	Quantum Monte Carlo Modeling of Strongly Correlated Electronic Systems	Physics	2,500,000
QMCPACK_aesp	Anouar Benali	Argonne National Laboratory	Extending Moore's Law Computing with Quantum Monte Carlo	Materials Science	30,000,000
QMCwater	Dario Alfe	University College London (UCL)	Water at Surfaces from Highly Accurate Quantum Monte Carlo Calculations	Materials Science	2,000,000
QMC_FeComplex	Hanning Chen	The George Washington University	Quantum Monte Carlo Study of Spin-Crossover Transition in Fe(II)-based Complexes	Materials Science	7,000,000
QuantumDS	Alvaro Vazquez-Mayagoitia	Argonne National Laboratory	Quantum Mechanics and Data Science	Chemistry	1,000,000
radix-io	Philip Carns	Argonne National Laboratory	System Software to Enable Data-Intensive Science	Computer Science	2,000,000
SC17SCC	William Scullin	Illinois Institute of Technology (IIT)	IIT Student Cluster Challenge Preparation: 2017	Computer Science	500,000
Slip-Flows	Ratnesh Shukla	Indian Institute of Science Bangalore	DNS of Flow Past Hydrodynamically Slipping Surfaces	Engineering	1,500,000
SolarWindowsADSP	Jacqueline M. Cole	University of Cambridge	Data-Driven Molecular Engineering of Solar Powered Windows	Materials Science	54,000,000
SU2_PadeOps_aesp	Sanjiva Lele	Stanford University	Benchmark Simulations of Shock-Variable Density Turbulence and Shock-Boundary Layer Interactions with Applications to Engineering Modeling	Engineering	5,000,000
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	500,000
UINTAH_aesp	Martin Berzins, John Schmidt	University of Utah	Design and Evaluation of High-Efficiency Boilers for Energy Production Using a Hierarchical V/UQ Approach	Chemistry	7,000,000
UnsAdaptCFDPrio_tesp	Kenneth Jansen	University of Colorado-Boulder	Extreme Scale Unstructured Adaptive CFD: From Multiphase Flow to Aerodynamic Flow Control	Engineering	10,000,000
UnsAdaptCFD_tesp	Kenneth Jansen	University of Colorado-Boulder	Extreme Scale Unstructured Adaptive CFD: From Multiphase Flow to Aerodynamic Flow Control	Engineering	10,000,000
visualization	Michael E. Papka	Argonne National Laboratory	Visualization and Analysis Research and Development for ALCF	Internal	2,000,000
Viz_Support	Joseph Insley	Argonne National Laboratory	Visualization Support	Computer Science	2,000,000
WaterHammer	Hong Zhang	Argonne National Laboratory	Water Hammer Simulation	Mathematics	10,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
WindTurb_tesp	Juan Alonso, Sanjiva K. Lele, Thomas D. Economon, Ramesh Balakrishnan	Stanford University	Scale-Resolving Simulations of Wind Turbines with SU2	Energy Technologies	70,000,000
XGC_aesp	C.S. Chang	Princeton Plasma Physics Laboratory (PPPL)	High Fidelity Simulation of Fusion Reactor Boundary Plasmas	Fusion Energy	3,000,000
				Total Theta DD	1,231,173,028

About Argonne National Laboratory
Argonne is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC under contract DE-AC02-06CH11357. The Laboratory's main facility is outside Chicago, at 9700 South Cass Avenue, Lemont, Illinois 60439. For information about Argonne and its pioneering science and technology programs, see www.anl.gov.

DOCUMENT AVAILABILITY

Online Access: U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free at OSTI.GOV (http://www.osti.gov/), a service of the US Dept. of Energy's Office of Scientific and Technical Information.

Reports not in digital format may be purchased by the public from the National Technical Information Service (NTIS):

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Road
Alexandria, VA 22312
www.ntis.gov
Phone: (800) 553-NTIS (6847) or (703) 605-6000
Fax: (703) 605-6900
Email: **orders@ntis.gov**

Reports not in digital format are available to DOE and DOE contractors from the Office of Scientific and Technical Information (OSTI):

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
www.osti.gov
Phone: (865) 576-8401
Fax: (865) 576-5728
Email: reports@osti.gov

Argonne Leadership Computing Facility
Argonne National Laboratory
9700 South Cass Avenue, Bldg. 240
Lemont, IL 60439-4832

www.anl.gov

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.