ANL/ALCF-20/05

Argonne National Laboratory

Argonne Leadership Computing Facility

2018 Operational Assessment Report

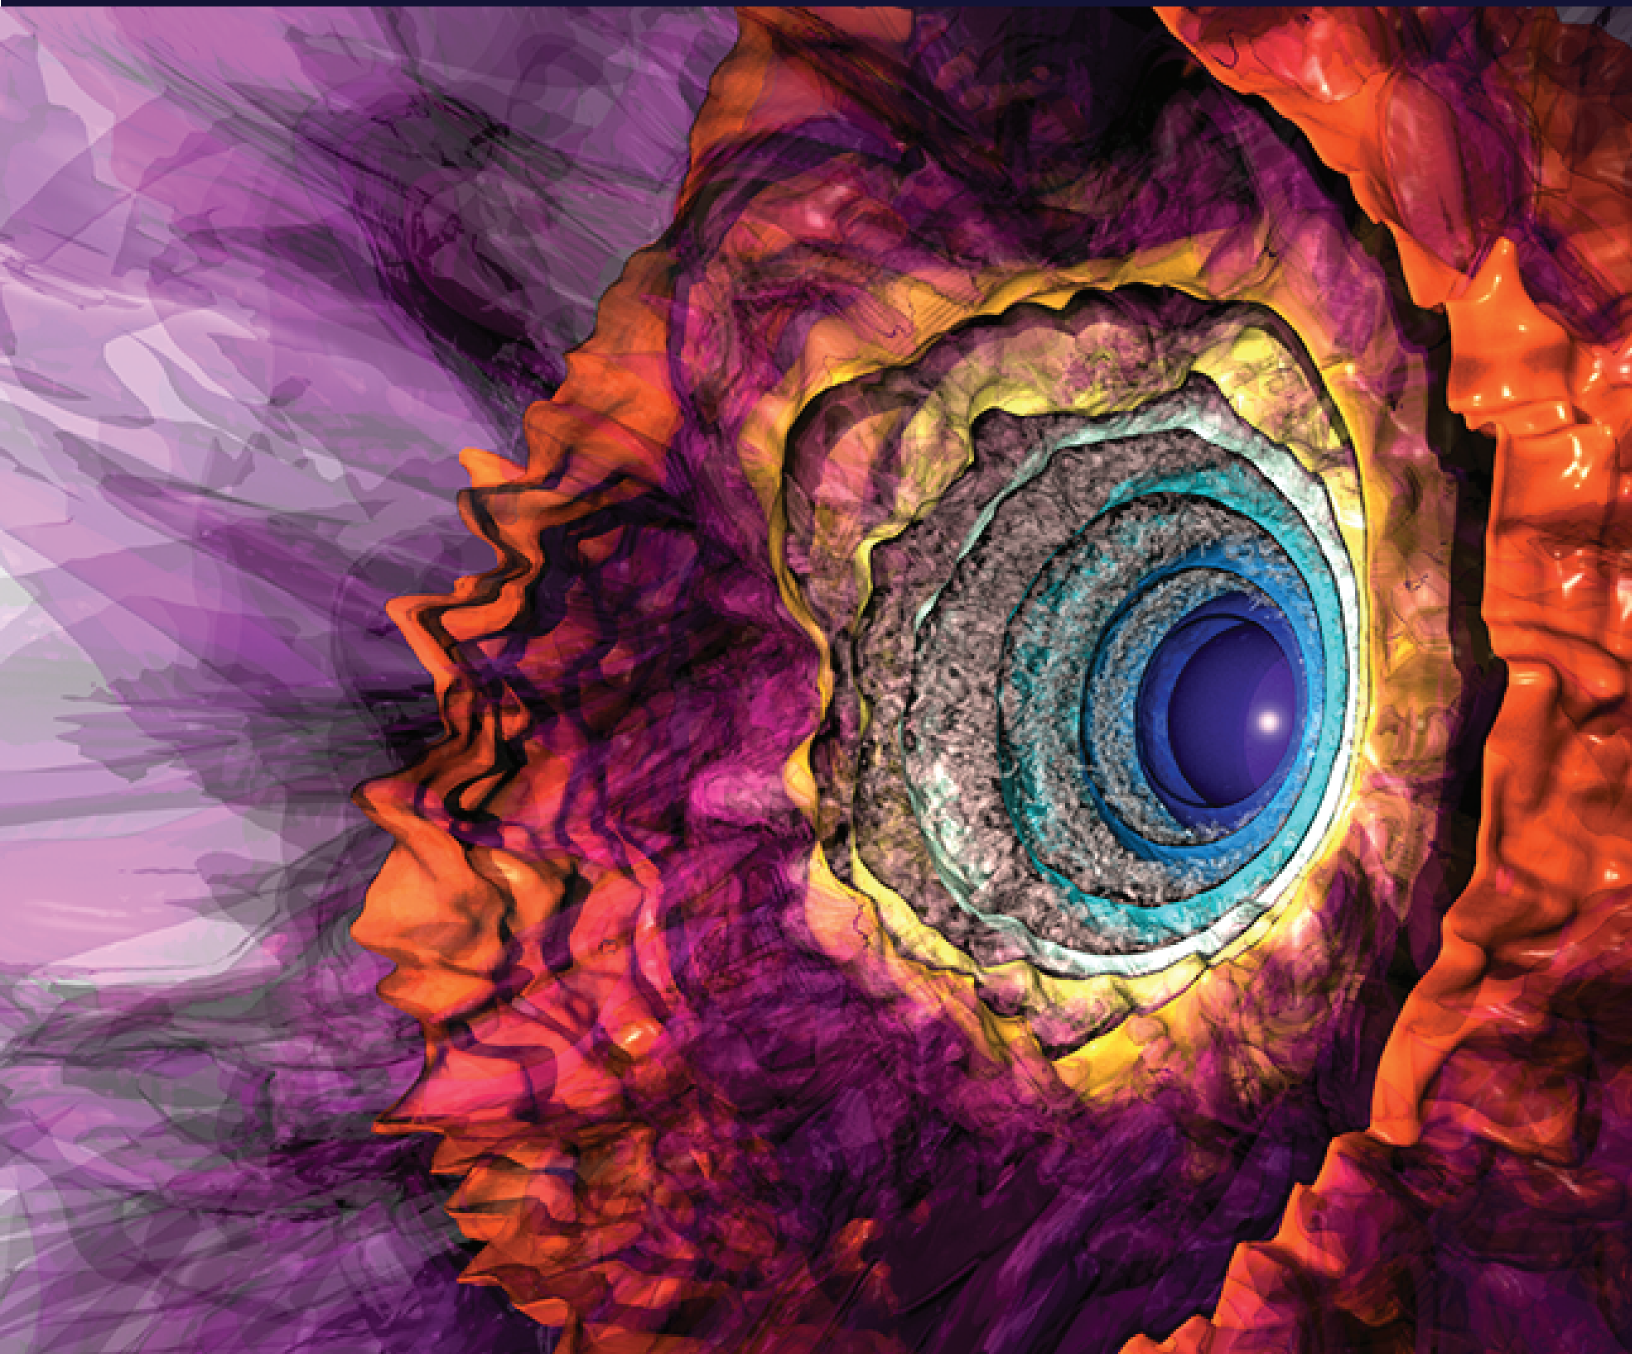

***On the cover**: A visualization of the turbulent envelope of a luminous blue variable star surrounding the central high-density core. These massive, evolved stars display large variations in their luminosity. Using Argonne National Laboratory's Mira supercomputer, Yan-Fei Jiang and his colleagues have found that helium opacity plays an important part in triggering the outbursts and in setting their effective temperature of around 9,000 Kelvin. The team predicts that these stars should also show variability at the 10–30% level on a timescale of days. The team's research was published in 'Nature' and also featured on the journal's cover. [DOI:10.1038/s41586-018-0525-0]*

***Visualization:** Joseph A. Insley/Argonne National Laboratory. **Science team:** Yan-Fei Jiang, Lars Bildsten, Omer Blaes, University of California, Santa Barbara; Matteo Cantiello, Flatiron Institute; Eliot Quataert, University of California, Berkeley; James Stone, Princeton University.*

Computing time for this work was awarded through the U.S. Department of Energy's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program.

Contents

Contents (Cont.)

Contents (Cont.)

Contents (Cont.)

Contents (Cont.)

Figures

Figures (Cont.)

Figures (Cont.)

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 (Argonne) and at Oak Ridge National Laboratory (Oak Ridge) 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 performed 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 for 2018 as an advanced scientific computing center.

The ALCF operates two production petascale resources for the research community: An Intel-based Cray system, Theta, and its IBM system, Mira.

In 2018, Theta delivered 529.4 million core-hours to 14 Innovative and Novel Computational Impact on Theory and Experiment (INCITE) projects and 508.2 million core-hours to ASCR Leadership Computing Challenge (ALCC) projects (10 awards for the 2017–2018 ALCC year and 6 awards for the 2018–2019 ALCC year), as well as substantial support to Director's Discretionary (DD) projects (916.6 million core hours). As Table ES.1 shows, Theta had an excellent first full year of operation in terms of overall availability (94.3 percent), scheduled availability (98.0 percent), and utilization (92.1 percent; Table 2.1).

Meanwhile, Mira delivered 4.1 billion core-hours to 27 INCITE projects and 1.7 billion core-hours to ALCC projects (22 awards for the 2017–2018 ALCC year and 16 awards for the 2018–2019 ALCC year), as well as significant support to a wide range of DD projects (472.2 million core-hours). Yet again, as Table ES.1 shows, Mira had an excellent year in terms of overall availability (96.5 percent), scheduled availability (99.3 percent), and utilization (93.6 percent; Table 2.1).

Moreover, as of the submission of this document, the ALCF's user community has published 276 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 Previous Year (2018) Metrics

Area	Metric	2018 Target	2018 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	4.5/5.0
	User Survey – **Response Rate**	25.0%	46.9%
	% User Problems Addressed within Three Working Days	80.0%	96.1%
Business Results	Theta Overall Availability	80.0%	94.3%
	Theta Scheduled Availability	90.0%	98.0%
	Mira Overall Availability	90.0%	96.5%
	Mira Scheduled Availability	90.0%	99.3%
	% of INCITE core hours from jobs run on 20.0% or more of Theta (51,200–281,088 cores)	20.0%	64.4%
	% of INCITE core hours from jobs run on 60.0% or more of Theta (153,600–281,088 cores)	N/A [a]	13.3%
	% of INCITE core hours from jobs run on 16.7% or more of Mira (131,072–786,432 cores)	40.0%	70.8%
	% of INCITE core hours from jobs run on 33.3% or more of Mira (262,144–786,432 cores)	10.0%	22.0%

[a] N/A = not applicable.

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 2018 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). User satisfaction with ALCF services remain consistently high as evidenced by survey response data. 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) 2018.

Table 1.1 All 2018 User Support Metrics and Results [a]

		2017 Actual	2018 Target	2018 Actual
Number Surveyed		976	N/A [b]	954
Number of Respondents (Response Rate)		458 (46.9%)	25.0%	447 (46.9%)
Overall Satisfaction	Mean	4.5	3.5	4.5
	Variance	0.5	N/A	0.6
	Standard Deviation	0.7	N/A	0.8
Problem Resolution	Mean	4.6	3.5	4.5
	Variance	0.5	N/A	0.5
	Standard Deviation	0.7	N/A	0.7
User Support	Mean	4.5	3.5	4.5
	Variance	0.5	N/A	0.5
	Standard Deviation	0.7	N/A	0.7
		2017 Actual	**2018 Target**	**2018 Actual**
% of Problems Addressed Within Three Working Days [c]		95.8%	80.0%	96.1%

[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] N/A = not applicable.

[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

In 2017, the ALCF worked with a consultancy, Marketing Synergy, Inc. (MSI), to reduce the overall length of its annual survey. This aim was achieved primarily by removing workshop-related questions and retaining the questions that comprised the DOE metrics for the 2017 Operational Assessment Report (OAR). The facility now conducts separate targeted surveys to capture user sentiments about ALCF-sponsored workshops. The results of the annual user survey dating back to 2008 appear on the ALCF website.

The 2018 user survey was administered by MSI and contained 24 questions. The survey was similar to the 2017 user survey, with two additional inquiries to capture sentiments about aspects of the ALCF's user services. Most respondents were able to complete the survey in 10 minutes or less. The 2018 survey e-mail campaign commenced November 22, 2018, and ended in early December 2018, once the minimum response rate (25.0%) was reached. Each reminder e-mail included 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-point scale. This is a standard way to rate user responses for surveys because (1) it provides a symmetric agree-disagree scale; (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 also allows for use of off-the-shelf statistics functions to determine variance and standard deviation.

The 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 [a]	(No Value)

[a] N/A = not applicable.

Only the Overall Satisfaction question applied a different point scale, as follows:

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

Beginning in 2017, some of the non-metric survey questions were revised to capture sentiments about various aspects of the 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

In CY 2018, 954 individuals met the definition of a facility user and were invited to complete the annual user survey. Of this group, 447 responded, for a 46.9 percent response rate — far superior to a generally accepted standard for survey response rates of 10 percent for this population size. The 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 higher average Overall Satisfaction than ASCR Leadership Computing Challenge (ALCC) and Director's Discretionary (DD) users, the results are not statistically significant. Other metrics are comparable, in that the variations are statistically insignificant. In their analysis, MSI noted that the historic overall user support ratings are extremely high compared to similar industry surveys.

Table 1.2 2018 User Survey Results by Allocation Program

2018 Metrics by Program		INCITE	ALCC	INCITE + ALCC	DD	All
Number Surveyed		316	256	572	382	954
Number of Respondents		159	141	300	147	447
Response Rate		50.3%	55.1%	52.4%	38.5%	46.9%
Overall Satisfaction	Mean	4.6	4.4	4.5	4.4	4.5
	Variance	0.5	0.5	0.5	0.8	0.6
	Standard Deviation	0.7	0.7	0.7	0.9	0.8
Problem Resolution	Mean	4.6	4.5	4.6	4.5	4.5
	Variance	0.4	0.5	0.4	0.5	0.5
	Standard Deviation	0.7	0.7	0.7	0.7	0.7
User Support	Mean	4.6	4.5	4.5	4.5	4.5
	Variance	0.5	0.5	0.5	0.5	0.5
	Standard Deviation	0.7	0.7	0.7	0.7	0.7
All Questions	Mean	4.6	4.5	4.5	4.5	4.5
	Variance	0.5	0.5	0.5	0.5	0.5
	Standard Deviation	0.7	0.7	0.7	0.7	0.7

As Table 1.3 shows, in 2018, the ALCF again exceeded the targets for overall satisfaction and user support.

Table 1.3 2017 and 2018 User Support Metrics

Survey Area	2017 Target	2017 Actual	2018 Target	2018 Actual
Overall Satisfaction Rating	3.5/5.0	4.5/5.0	3.5/5.0	4.5/5.0
Avg. 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

	2017 Target	2017 Actual	2018 Target	2018 Actual
% of Problems Addressed Within Three Working Days [a]	80.0%	95.8%	80.0%	96.1%
Avg. of Problem Resolution Ratings	3.5/5.0	4.6/5.0	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

In CY 2018, the ALCF answered and categorized 7,172 support tickets. All 20 categories came within one percentage point of last year's numbers, with the following six exceptions: 'Access' went from 15 percent in 2017 to 18 percent in 2018; 'Accounts' went from 46 percent in 2017 to 41 percent in 2018; 'File System' and 'HPSS and Quota Management' went from 1 percent in 2017 to 3 percent in 2018; 'Scheduling' went from 2 percent in 2017 to 4 percent in 2018; and 'Reports' went from 1 percent in 2017 to 4 percent in 2018. Once again, the 'Accounts' category netted the largest number of tickets overall (see Table 1.5).

The ALCF is planning to distinguish between incidents that are problems versus routine requests when we migrate to a different ticketing software system in the future.

Table 1.5 Ticket Categorization for 2017 and 2018

Category	2017	2018 Total
Access	923 (15%)	1,323 (18%)
Accounts	2,870 (46%)	2,948 (41%)
Allocations	594 (10%)	621 (9%)
Applications Software	176 (3%)	236 (3%)
Automated E-mail Responses	560 (9%)	640 (9%)
Compilers	52 (1%)	73 (1%)
Data Transfer	38 (1%)	37 (1%)
Debugging and Debuggers	12 (0%)	37 (1%)
File System	81 (1%)	221 (3%)
HPSS [a] and Quota Management	60 (1%)	211 (3%)
Libraries	18 (0%)	16 (0%)
I/O and Storage [b]	103 (2%)	-
Miscellaneous	197 (3%)	167 (2%)
Network	6 (0%)	21 (0%)
Performance and Performance Tools	12 (0%)	46 (1%)
Reports	70 (1%)	289 (4%)
Quota Management [c]	25 (0%)	-
Scheduling	107 (2%)	267 (4%)
System [d]	295 (5%)	-
Visualization	15 (0%)	19 (0%)
TOTAL TICKETS	6,214 (100%)	7, 172 (100%)

[a] HPSS = high-performance storage system.

[b] The I/O (input/output) and Storage category was removed at the end of 2017 2Q and replaced with HPPS and Quota Management.

[c] Quota Management was removed at the end of 2017 2Q and replaced with File System.

[d] System was replaced by seven new categories (‘Compilers,’ ‘Debugging and Debuggers,’ ‘Libraries,’ ‘Network,’ ‘Performance and Performance Tools,’ ‘Reports,’ and ‘Scheduling’) from 2017 3Q.

1.3.1.2 Improved Account Renewal Communication

The automated Account Renewal e-mails sent to users now include a username in the subject line. This change allows account sponsors/principal investigators (PIs) to quickly scan multiple pending renewal requests without first having to open each e-mail.

1.3.1.3 Process Optimization for Allocation Ramp-ups

A process improvement for setting up allocation awards on ALCF systems was implemented by partially automating the allocation creation and file space assignments for approved projects. This process was used to set up 2018–2019 ALCC awards, 2018–2019 ALCF Data Science Program (ADSP) awards, Data and Learning Early Science Program (ESP) awards, INCITE 2019 awards, and numerous DD awards.

1.3.1.4 Integration and Extension of Banhammer

Until 2018, the ALCF used a stand-alone Banhammer system to assign penalties on Mira for ALCC, ADSP, and INCITE projects that failed to meet quarterly reporting requirements. In early 2018, this feature was integrated with sbank, the ALCF's allocation accounting system, and extended to Theta and Cooley.

1.3.1.5 Streamlined Quarterly Reports Collection Workflow

The collaborative effort begun last year to improve the workflow for collecting quarterly and annual reports was further streamlined and standardized for all of the allocation programs. The internal projects checklist was migrated from MediaWiki to the ALCF's document repository system, Confluence, which simplified the process of editing and updating a project's report status. In addition, the report templates were converged to produce one unified set for all of the allocation programs.

1.3.2 Application Support

1.3.2.1 Individual Projects

Interfacial Behavior of Alcohol at Water/Organic Biphasic System

For the DD project "WOInterface — Interfacial Behavior of Alcohol at Water/Organic Biphasic System," ALCF staff provided comprehensive support including NAMD code porting/usage, parameter preparation, and simulation design. This project resulted in a publication in *The Journal of Physical Chemistry C*. The project also provided the first large-scale benchmark for liquid/liquid separation simulation on a DOE leadership supercomputer and resulted in an ongoing Laboratory Directed Research and Development (LDRD) project (PRJ1006500 Complex Fluids from Chemical Separations).

Optimization and Performance Data for INCITE Submission

The goal of this DD project was to improve the strong scaling of VSVB energy calculations on the iron molybdenum cofactor (FeMo-co) of nitrogenase. FeMo-co energy calculations have low granularity due to expensive density factors (many openshell, singlet-coupled electrons), resulting in poor load balancing. Strong scaling is poor unless the task count is an exact multiple of the number of ranks; however, the task count is unknown in advance because of integral screening. The opportunity arises that each density factor is a double-sum over determinants, in two nested loops. Thus, a possible strategy was to subdivide the density computation inside one loop, creating a longer list of smaller tasks that can be distributed over an arbitrary number of ranks without incurring large load imbalances. Implementing this strategy improved the parallel efficiency from 31 percent to 85 percent on 65,536 cores of Theta.

A Systematic Study of Minima in Alanine Dipeptide

ALCF computational science support staff and David Baker's group at the University of Washington performed a systematic study of minima in alanine dipeptide. The alanine dipeptide is a standard system to model dihedral angles in proteins. Obtaining the Ramachandran plot accurately is known as a difficult problem because of many local minima; depending on the details of geometry optimizations, different Ramachandran plots can be obtained. To locate all energy minima, starting from geometries from molecular dynamics (MD) simulations, 250,000 geometry optimizations were performed at the level of RHF/6-31G*, followed by re-optimizations of the located 827 minima at the level of MP2/6-311++G**. This treatment yielded 30 unique minima, most of which were not previously reported in the literature. Both *in vacuo* and solvated structures were discussed in the findings. The minima are systematically categorized based on four backbone dihedral angles. The Gibbs energies were evaluated, and the structural factors determining the relative stabilities of conformers were discussed. The results of this work were accepted into the special issue of *Journal of Computational Chemistry*. This work was also highlighted on the journal's cover.

Access to Pre-release Version of VASP 6

In support of various DD, INCITE, and ADSP projects that use Theta, and in collaboration with Intel, an ALCF staff member obtained access to a pre-release version of VASP 6. This version of VASP is highly optimized for the Knights Landing (KNL) architecture over the latest release of VASP 5.4.x. The pre-release version of VASP 6 is up to 3× faster for hybrid functionals over VASP 5.4.x. This version of VASP was optimized by the University of Vienna and Intel. The ALCF obtained privileged access to VASP 6 and compiled it on Theta.

Investigation of a Low Octane Gasoline Fuel for a Heavy-Duty Diesel Engine

Researchers from Argonne National Laboratory (Argonne) and Aramco Services Company (ASC) – Detroit are using ALCF resources to accelerate the optimization of engine design. For one study, Argonne's Virtual Engine Research Institute and Fuels Initiative (VERIFI) team used Mira to design and optimize the combustion chamber of a heavy-duty diesel engine using a gasoline-like fuel for improved efficiency and performance. This first-of-its kind study used a high-fidelity simulation approach to optimize the fuel spray and combustion bowl geometry on a supercomputer. The team simulated more than 2,000 engine design combinations and reduced design time from months to weeks for ASC. Simulations demonstrated the potential for significant fuel efficiency benefits. An ALCF performance engineer helped optimize Convergent Science's CONVERGE code on Mira, resulting in a 100× speedup in I/O, 8× in load balance, and 3.4× in time-to-solution, among other improvements. An ALCF staff member assisted with the visualizations.

Lithium-Oxygen Battery with Long Cycle Life in a Realistic Air Atmosphere

This 2017 DD project on Vesta led by University of Illinois at Chicago and Argonne produced a new design for a Li-O_2 battery cell that operates by a reaction with air over many charge and discharge cycles. Advances in materials functionality for Li-O_2 electrochemistry have resulted in a Li-O_2 battery that is able to run under a realistic air atmosphere with a long cycle life. The new Li-O_2 cell architecture is a promising step toward engineering the next generation of lithium batteries with much higher specific energy density than lithium-ion batteries. The team's basic science computational studies revealed how this system operates in air and what factors contribute to the improved cycling stability: Specifically, the team performed voltage profiles for a Li-O_2 cell based on a molybdenum disulfide (MoS_2) cathode, ionic liquid/dimethyl sulfoxide electrolyte, and a lithium carbonate-coated lithium anode with a cycle life of more than 500 cycles. Results of an *ab initio* molecular dynamics simulation using VASP of the electrolyte used in the Li-O_2 cell helped explain its stability. The results of this study were reported in the journal *Nature*.

Balsam: Workflow Manager and Edge Service for ALCF Systems

Using DD awards on Mira and Theta, the ALCF's Data Science team has developed the Balsam workflow manager to simplify the task of running large-scale job campaigns on ALCF resources. The DIII-D National Fusion Facility used Balsam to trigger experiment-time analyses during Tokamak operation, running more complex analyses in less time and leading to higher experiment productivity. The ATLAS High Energy Physics (HEP) experiment at CERN has used Balsam to run hundreds of millions of compute hours of event generation jobs on ALCF systems. ATLAS HEP has used Balsam for production science on the National Energy Research Scientific Computing Center (NERSC) system as well, demonstrating the portability and efficacy of Balsam. The ALCF is currently using Balsam to conduct hyperparameter optimization for deep learning at scale on Theta. Other potential projects include computation for real-time Advanced Photon Source (APS) experiments, ADSP, and Exascale Computing Project (ECP) projects.

Scaling ATLAS Simulation Workflows

This project modified the High Throughput Computing model used by ATLAS for the past 10 years on the Large Hadron Collider (LHC) grid to run on Theta. Working with ALCF staff, Argonne physicists developed a message passing interface (MPI)-enabled workflow to run the Geant4 simulation for production science. Containers for software deployment on Theta were crucial to scaling and running the multistep simulation workflow. Combined with a new data streaming framework developed specifically for supercomputers, 65M simulated events have been delivered to ATLAS using 50M core-hours on Theta since September 2017. In addition, 2.5M events were generated for studying machine learning algorithms for particle reconstruction in a three-dimensional (3D) detector environment.

Modeling Electronic Stopping in Condensed Matter under Ion Irradiation

Researchers from the University of North Carolina at Chapel Hill and the University of Illinois at Urbana-Champaign are using their predictive framework for electronic stopping to model scenarios that could potentially lead to advancements in such applications as proton-beam cancer therapy. Using magnesium oxide for the target and silicon for the projectile, the researchers have completed their investigation of the initial-condition dependency of electronic stopping for heavy ions and the critical role played by (semi-)core electrons. This dependency had never been observed before; and, while an unfavorable feature, it was discovered that long-duration simulations on leadership-class supercomputers can provide its resolution. The team now seeks to advance the first-principles electronic stopping capabilities for heavy ions in semiconductors and to explore the intricacies of irradiating deoxyribonucleic acid (DNA) with light ions like protons, alpha particles, and carbon ions. A 2018 INCITE project award on Mira and Theta allowed the team to continue to compare proton and alpha-particle irradiation of solvated DNA. Calculations quantifying the presence of holes in DNA and the charge removed by projectile ions provided further molecular-level insights. Simulations of hydrated DNA are ongoing. ALCF staff provided basic support and assistance with using additional runtime beyond the original award.

Accelerated Climate Modeling for Energy

The Accelerated Climate Modeling for Energy (ACME) project, now known as the Energy Exascale Earth System Model, or E3SM, aims to develop a cutting-edge climate and Earth system model that can tackle the most demanding climate research imperatives. A Sandia National Laboratories-led research team is focused on the water cycle and cryosphere systems using resources at the ALCF and Oak Ridge Leadership Computing Facility (OLCF). The team is first seeking to simulate changes in the hydrological cycle, specifically focusing on precipitation and surface water in regions where this cycle is impacted by complex topography. The second objective is to determine the possibility of dynamical instability in the Antarctic Ice Sheet during the next 40 years. The team made extensive use of Theta to run the recently released E3SM code, which comprises component models for atmosphere, ocean, sea ice, and land. The efficient execution of each simulation involves an optimal balance of compute nodes among the various component models, with the majority of the nodes being allocated to the atmosphere model, a subset of which also runs the land and sea models. The remaining nodes are allocated to the ocean model, which runs concurrently with the atmosphere model. The examination of the risk of Antarctic Ice Sheet collapse represents the first fully coupled simulation to include dynamic ocean-ice shelf interactions. The team also conducted several simulations ranging from 3–5 years in length to test different atmospheric tunings and initial conditions for ocean and ice. Beyond evaluation of the model, this testing also allowed for workflow development and prepared the team for deeper studies. ALCF staff coordinated with the project team to provide additional runtime beyond the original award of 179 million core-hours split between the ALCF and OLCF.

1.3.3 Resource Support

1.3.3.1 Cooley RHEL 7 Upgrade

In fiscal year (FY) 2018, we upgraded Cooley's operating system to Red Hat Enterprise Linux (RHEL) 7. Upgrading to the current release of RHEL allows us to provide our users with the most recent software packages available from Red Hat and EPEL and, in turn, allows users to take advantage of new software features and bug fixes, and reduces the need for users to compile software directly from the source code from third-party sources to obtain more current versions. In addition, RHEL 7 has first priority for security updates, which reduces the length of exposure when operating system security vulnerabilities are discovered.

1.3.3.2 Mounting Lustre on Cooley

As a result of the Cooley RHEL 7 upgrade, ALCF staff members were also able to install a supported version of the Lustre client on Cooley to make Theta's file systems directly available to users. This change allows users to perform data analytics and visualization of Theta output on Cooley without any need to transfer the data using Globus or other file transfer mechanisms. This structure also reduces total storage usage, as it eliminates the need to copy data from Lustre to general parallel file systems (GPFSs).

1.3.3.3 LTO8 Upgrade

A tape technology upgrade in July 2018 provided a variety of benefits to the tape archive infrastructure. With a native capacity of 12 terabytes (TB), the LTO8 tape media supplies 5 times the capacity of the previous media, LTO6. In addition, the transfer rate of 360 megabytes (MB)/s achieved by the LTO8 tape drive is more than twice that of the LTO6 drive. The data density provided by the increased capacity will allow, over time, hundreds of tape library slots to be freed and made available for new tape to provide additional storage capacity. This process, known as repacking, will be ongoing until data are migrated off the LTO6 media.

1.3.3.4 Cobalt Updates

Work continues on the Cobalt scheduler and resource manager to better support our users and new workloads. Users are now able to request the on-node solid state drives (SSD) available on Theta. This capability also ensures that all nodes assigned to a job have sufficient SSD storage available for that job. Built-in support for cgroup isolation of user scripts has also been made available, which enables us to better support users that may put higher demands on shared login and job monitoring resources. In addition, we have made numerous improvements to the scheduling and stability of Cobalt on Cray systems as we have diversified the users and workloads on Theta over the last year, while maintaining stability and high utilization on all ALCF resources.

We are also planning a major design revision to Cobalt to scale to exascale and anticipated post-exascale architectures, and to better support highly diverse workloads and novel hardware and capabilities. We have upgraded Cobalt's testing to incorporate our GitLab and Jenkins Continuous Integration (CI) infrastructure to not only improve test coverage but also to automate testing against all supported operating system (OS) and system environments. (Documentation of the ALCF's CI efforts appear in Section 4 and Section 8.) This approach, combined with significant improvements in logging and messaging from Cobalt, allows us to confidently support our systems running Cobalt while work proceeds on this revision.

1.3.3.5 Ram Area Network (RAN)

The ALCF and Chicago-based tech company Kove IO, Inc., have been running a multiyear background research project to look at memory disaggregation in high-performance computing (HPC). The ALCF Cooley visualization cluster has access to 6 TB of pooled random access memory (RAM) via 14 FDR InfiniBand (IB) cards installed across three appliances.

Disaggregated memory has a number of potentially significant impacts, which include the following:

- Reduction of aggregate system RAM and its associated purchase and operations costs.
- The ability to have virtually any amount of RAM on a node, on demand.
- The decoupling (mostly) of purchases of central processing units (CPUs), which we need to upgrade every few years to obtain improved processor performance, from RAM, whose performance has changed very little over the past decade.
- The ability to eliminate guessing at the number of "big memory nodes" a heterogeneous cluster needs and having its RAM sit idle, consuming power, when not in use.

The obvious issue is latency and its impact on performance. ALCF-supported student researchers have tried a number of different application areas to see how they behave. An Illinois Institute of Technology student is currently following up on early results suggesting that machine learning and deep learning are excellent use cases for disaggregated memory. This student is greatly expanding the testing regime to understand what characteristics of the algorithms make it so amenable. The student has expressed her intention to pursue a Ph.D. focused on disaggregated architectures, so we expect to be working with her for several years to come.

ALCF staff members will continue to explore how disaggregated memory fits into the HPC deep memory hierarchy. In 2018, we had several students working on a variety of topics including:

- Extending the Concurrent Average Memory Access Time (C-AMAT) model to include disaggregated memory. In particular, during some of our testing we found that a K-Means benchmark performed better with RAN than with local memory, which was a surprising result. Using the C-AMAT model, we believe this result is attributable to better locality in the L1D cache due to concurrent access by multiple threads.
- Beginning development of simulator modules to simulate disaggregated memory.
- Conducting further testing that continues to support deep learning as an excellent candidate for disaggregated memory. Results continue to show that with a cache of <=20% on the local machine, time-to-solution rates remain constant using disaggregated memory, enabling virtually unlimited training set sizes.
- Developing tools to monitor and plot performance metrics related to disaggregated memory.
- Developing a prototype for using disaggregated memory for buffering during two-phase I/O.

1.3.3.6 User Management System

The application used to manage ALCF user data, called userbase, is a home-grown solution that was originally written many years ago within a different Argonne division and then adopted by the ALCF for its needs. While ALCF developers have maintained the application to ensure minimum security standards were met, the application itself lacked many features that staff members found they required as the ALCF evolved. Taking this opportunity to better define and

implement a solution, we decided to develop a new application to manage users. Most of the development work on this new application was completed in 2018, and we are now in a beta testing phase with planned deployment in the first quarter of 2019.
In creating this new application, here are some best practices and improvements that were used:

- Integration with the Argonne Foreign Assignment/Visit Request (FAVOR) system. The old version of userbase relied on manual, often error-prone, transfers of data between the ALCF and the lab-wide system. This also meant the old application stored personally identifiable information (PII). The new application sends the users directly to a form within the lab-wide system so as to by-pass the entry of PII into the ALCF application altogether. Using an application programming interface (API), relevant data that the ALCF needs are brought back to display within the application. Argonne's Business and Information Services (BIS) FAVOR team is recommending this methodology to other divisions at Argonne. We will capture our process and lessons learned so we can help other divisions leverage this work as needed.
- Currently, different computing groups within the Computing, Environment and Life Sciences (CELS) directorate run their own, forked versions of userbase. The new application is written to have a centralized, shared identity database and then to support an arbitrary number of "domains." The domains can all be executed on a centralized host or each domain can run its own instances; but in all cases, the identity information is shared for easy access between computing groups at Argonne.
- A number of software development best practices have been put in place including:
  - All functionality is available via RESTful APIs.
  - We developed a tool that scans the back-end Python code and automatically generates APIs to be called by the front-end AngularJS code.
  - We use Vagrant to have a consistent and portable environment from our development on laptops through test, stage, and production servers.
  - The front end and back end can be independently developed and tested with mocks for the other side prior to integration testing.
  - We use table-driven data validation, which guarantees that the front end and back end are validating consistently.
  - Role-based security authentication and authorization are used for every RESTful API, which allows more granular security control.
  - Two-factor authentication is used for all users.
  - Modern design and branding are consistent with the ALCF website.
  - We make use of Gitlab for cross-team collaborations on development life cycles.
  - We make use of an internally hosted Jenkins Automation Server for CI.

1.3.4 Outreach Efforts

1.3.4.1 General Outreach

Facility tours are frequent and welcome opportunities for ALCF staff members to interact with the public and invited guests about ALCF research and activities. A tour coordinator vets all tour requests, both internal and external, and locates the appropriate staff member or members to conduct the tour. In 2018, the ALCF hosted approximately 75 groups.

User Advisory Council

The User Advisory Council (UAC) meets on a quarterly basis to offer recommendations for program improvements and to provide information relevant to policies and practices that affect the user community. The members represent a cross section of active community users and stakeholders served by the facility. UAC meetings resumed in August after a several-month break, during which a new liaison was appointed, two members left, four appointments were renewed, and two members were added. In two quarterly meetings held in August and November, the council provided feedback on the ALCF scheduling process and discussed the 2018 survey.

Connection to Technology Commercialization and Partnerships Division

The ALCF interacts with Argonne's Technology Commercialization and Partnerships (TCP) division and regularly supports TCP-organized meetings with potential industry partners. TCP in turn provides potential partners, even those without immediate HPC needs, with a complete picture of Argonne computing capabilities and partnership opportunities, including those at the ALCF.

1.3.4.2 Workshops, Webinars, and Training Programs

The ALCF regularly conducts workshops, webinars, and training programs to support the efforts of its users. The facility also collaborates with peer institutions to develop training opportunities, explore key technologies, and share best practices that improve the user experience.

Getting Started Videoconferences

The ALCF conducts Getting Started videoconferences for new users that can also serve as refresher training. The sessions cover system specifications, code building, storage, OS and file systems, compilers, tools, queues, and other topics, and include hands-on exercises. The ALCF hosted one for Mira on January 24, 2018, and one for Theta/Cooley on January 26, 2018.

Computational Performance Workshop

This annual scaling workshop, held in May, is a cornerstone of the ALCF's user outreach program. The workshop attracts prospective INCITE users for talks and hands-on application tuning on Mira and Theta. Attendees work with ALCF staff, and tool and debugger vendors, with the goal of submitting an INCITE project proposal.

Simulation, Data, and Learning Workshops

The ALCF hosted two 3-day, hands-on workshops in 2018, one in February/March and one in October, to help users improve the performance and productivity of simulation, data science, and machine learning applications on ALCF systems. Users worked directly with ALCF staff and industry experts to learn how to use data science tools and frameworks at scale on ALCF systems. Topics covered included machine learning frameworks (TensorFlow, Horovod, PyTorch, etc.) and workflow management services using Balsam, notebooks, and containers to advance data science projects.

ALCF Developer Sessions

In 2018, the ALCF hosted 10 installments of the ALCF Developer Sessions for Theta users. Formerly named the ALCF Many-Core Developer Sessions, this live-presentation webinar series was created to spark discussions between developers and the early users of emerging high-performance computing hardware and software. Speakers in this series have included developers from Intel, Cray, Allinea (ARM), Paratools, Rice University, and the ALCF. In 2018, topics covered included workflow management, identification of memory bottlenecks, run-to-run performance variability, power efficiency, TensorFlow performance optimizations, tuning of Python workloads, and a variety of other performance profiling methodologies. Videos of these webinars will be available on the ALCF website soon.

ATPESC 2018

The ALCF ran its sixth successful Argonne Training Program on Extreme-Scale Computing (ATPESC) event from July 29–August 10, 2018. ATPESC is an intensive, two-week summer program organized around seven core program tracks and focused on HPC methodologies that are applicable to both current and future machine architectures, including exascale systems. A total of 73 graduate students, postdoctoral students, and early career researchers from more than 50 different institutions worldwide attended the program, which consisted of technical lectures, hands-on exercises, and dinner talks. To extend the reach of this program, 60 hours of video playlists of ATPESC presentations were uploaded to Argonne's YouTube training channel, and promotion is ongoing. ATPESC organizers work diligently to improve the program from year to year. In 2018, this effort included improvements in outreach, logistics, and management of the event with benefits for participants, lecturers, and support staff. New pages describing the Argonne tour and the ATPESC computing resources were included on the ATPESC website, enhancing the importance of both elements of the program. ATPESC is part of the Exascale Computing Project (ECP), a collaborative effort of the DOE Office of Science and the National Nuclear Security Administration.

1.3.4.3 Community Outreach

Youth Outreach

Women in STEM

Introduce a Girl to Engineering Day (IGED) is an annual lab-sponsored event in February that pairs approximately 100 local eighth graders with Argonne engineers and scientists for a day of presentations and hands-on activities focused on science, technology, engineering, and mathematics (STEM) careers. Event co-chair Liza Booker and many other ALCF staff members served as speakers, mentors, and activities supervisors at IGED 2018. In April, Argonne hosted the 31st Annual Science Careers in Search of Women (SCSW) Conference, a one-day event that provides high school girls with an opportunity to explore careers in STEM fields through a variety of interactions with Argonne staff. The students attended presentations and career panel discussions, facility tours, a poster session, and a career fair. ALCF staff members served in various roles at SCSW 2018.

In 2018, ALCF women also participated in Women in Science and Technology, Women in Statistics & Data Science, and the Grace Hopper Celebration of Women in Computing.

CodeGirls@Argonne Camp

In June, 25 seventh- and eighth-grade girls attended CodeGirls@Argonne, a two-day STEM camp that teaches the fundamentals of coding in the Python programming language. In addition to learning the basics of Python, the campers met and interviewed five women scientists who work at Argonne, including some from the ALCF. The group also toured the ALCF machine room and Visualization Lab.

Summer Coding Camp

The ALCF staff taught and mentored 30 local high school students at Argonne's Summer Coding Camp, July 16–20, 2018. The camp curriculum promotes problem-solving and teamwork skills through hands-on coding activities, such as coding with Python and programming a robot. The camp is a joint initiative of the ALCF and Argonne's Educational Programs Office.

Big Data Visualization Camp

Fourteen current and recent high school students attended Argonne's Big Data Visualization Camp, July 24–26, 2018, to learn how visualization helps researchers see patterns and relationships in big data that would be hard to see otherwise. The students programmed visualization tools in Python and worked with data from the Array of Things project, an Argonne–University of Chicago urban sensor project. The camp, a new STEM summer program that teaches computational thinking skills, is organized by Argonne's Educational Programs Office and taught by Argonne computer scientists.

Hour of Code

As part of the national Computer Science Education Week (CSEdWeek) in December, ALCF staff members visited various Chicago and suburban classrooms to give talks and demos and lead 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.

Conference Scholarship Program for Women in Technology

An ALCF staff member and member of the Portland Women in Technology (PDXWIT) operations team led the 2018 PDXWIT scholarship program for awarding small grants to women in Portland's tech community to attend technology-related or women-in-business conferences. The scholarship program was launched in June 2018 and announced its first awardees in December 2018. Four applicants were each awarded up to $2,500 to be applied to travel and registration.

Broader Engagement Activity at SIAM CSE

The Society of Industrial and Applied Mathematics (SIAM) 2019 Computational Science and Engineering (CSE19) conference's Broader Engagement program provided mentoring and career and professional development to students from underrepresented and underprivileged backgrounds who aspire to broaden their experience in research-based professional activities.

Within this program, an ALCF staff member managed the Guided Affinity Groups (GAGS), which paired students with a subject matter expert for the duration of the conference.

Visits and Tours

As noted in Sections 1.3.4.2 and 1.3.4.3, the laboratory hosts numerous school groups, government officials, and delegations throughout the year. Below is a selection of the groups and guests that toured the facility in 2018.

Government Officials

On March 7, 2018, DOE Under Secretary Mark W. Menezes toured the ALCF accompanied by Argonne Associate Laboratory Director Rick Stevens and other senior laboratory officials. Other tours included Suzanne Jaworowski, Chief of Staff and Senior Advisor in the DOE's Office of Nuclear Energy, on March 26, 2018, and DOE Office of the Under Secretary Senior Advisor Thomas Cubbage and Chief of Staff Kristin Ellis on April 5, 2018. The House Committee on Science visited on May 11, 2018, as did Tali Bar Shalom, White House Office of Management and Budget, and Kathleen Klausing, DOE Office of Science, on August 28, 2018. DOE Under Secretary for Nuclear Security Lisa E. Gordon-Hagerty toured the ALCF on August 30, 2018.

International Delegations

Alexander C. Kalloniatis, member of Australia's Defence Science and Technology Group, visited the ALCF on October 31, 2018. On November 3, 2018, ALCF staff provided a tour for ThorCon Power and delegates from the Indonesian Ministry of Energy and Mineral Resources.

Education Groups

On March 22, 2018, 50 honors students from Purdue University's College of Engineering toured the ALCF. In August, 15 members of the University of Wisconsin Research Systems Administration Group came for a day-long site visit to learn about and discuss potential areas for collaboration with Argonne. The group included staff from Project Ice Cube (neutrino physics), medical sciences, the Center for High Throughput Computing, social sciences, and biological sciences. On August 29, 2018, ALCF staff gave a tour as part of an International Atomic Energy Agency Training Workshop for International Middle and High School Science Teachers.

Industry Groups

Industry tours included representatives from the Oil & Natural Gas Sector Coordinating Council on June 20, 2018, and participants of the Exelon Microgrid Summit on September 13, 2018.

Other STEM Activities

The last week of April, an ALCF staff member was invited to speak at the 2018 DOE National Science Bowl finals in the last week of April, a nationwide academic competition for middle school and high school students that tests students' knowledge in all areas of science and mathematics. His talk covered a brief history of supercomputing and what the future might hold, including use of exascale machines.

In May, an ALCF staff member ran an improv activity on communicating science to children and parents as part of West Union Elementary STEAM (Science, Technology, Engineering, and Math + Art) Night in Portland, Oregon.

1.3.5 Communications

Communications through Mailing Lists and Social Media

The 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 2018 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 and text-format 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 2018 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 the 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, PIs/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 Argonne training activities, such as ATPESC, or fellowship opportunities, such as the Margaret Butler Fellowship in Computational Science.

Promotional Activities and Media Hits

In 2018, the ALCF produced and published 49 original science stories on the ALCF website and in its companion newsletter, Newsbytes. In coordination with Argonne's Communications and Public Affairs office and other ALCF direct relationships, ALCF science stories are disseminated to various news outlets (print, web, television, etc.). If the story passes through an editorial review and appears in a unique, individual media title, the item is counted and tracked as a media "hit." In 2018, the facility posted 93 such media hits to the website. The media team uses Meltwater News public relations suite to help track media hits. This global online media monitoring company tracks articles from more than 200,000 news publications, Twitter, YouTube, Facebook, and blogs. Meltwater reported 134 media hits and an audience reach of 47.8 million.

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 available on the ALCF website.

Table 1.8 Publications Designed for Print

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

Conclusion

Our users are at the forefront of all of the 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 2018, improvements included working with ALCC and INCITE project teams to transform and optimize their scientific codes and 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.

This page intentionally left blank.

Section 2. Operational Performance

Did the facility's operational performance meet established targets?

ALCF Response

The 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, the ALCF is on par with the other facilities and has demonstrated exceptional performance. To assist in meeting these objectives and to improve overall operations, the 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 Operational Performance Section

	Mira (Blue Gene/Q): 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				Theta (Cray XC40): 4,008-node, 257K-core, 64 TB MCDRAM,[a] 770 TB DDR4			
	CY 2017		CY 2018		CY 2017		CY 2018	
	Target	Actual	Target	Actual	Target	Actual	Target	Actual
Scheduled Availability	90.0%	99.5%	90.0%	99.3%	90.0%	98.6%	90.0%	98.0%
Overall Availability	90.0%	96.4%	90.0%	96.5%	80.0%	90.0%	90.0%	94.3%
System MTTI	N/A [a]	10.37 days	N/A	12.18 days	N/A	6.94 days	N/A	10.12 days
System MTTF	N/A	36.44 days	N/A	72.73 days	N/A	30.30 days	N/A	32.55 days
INCITE Usage	3.5B	4.5B [f]	3.5B	4.1B [h]	N/A	N/A	390.9M	529.4M
Total Usage	N/A	6.4B [g]	N/A	6.3B [i]	N/A	724.1M	N/A	2.0B
System Utilization	N/A	95.5%	N/A	93.6%	N/A	82.7%	N/A	92.1%
INCITE Overall Capability [b,d]	40.0%	72.2% [f]	40.0%	70.8% [h]	N/A	N/A	20.0%	64.4%
INCITE High Capability [c,e]	10.0%	27.5% [f]	10.0%	22.0% [h]	N/A	N/A	N/A	13.3%

[a] MCDRAM = multi-channel DRAM; DDR4 = Double Data Rate 4 Synchronous Dynamic Random-Access Memory; N/A = not applicable.

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

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

[d] Theta Overall Capability = Jobs using ≥ 20.0 percent (800 nodes, 51,200 cores) of Theta.

[e] Theta High Capability = Jobs using ≥ 60.0 percent (2400 nodes, 153,600 cores) of Theta.

[f] Usage includes 3604 core-hours from Cetus production jobs.

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

[h] Usage includes 0.3M core-hours from Cetus production jobs.

[i] Usage includes 11.1M core-hours from Cetus production jobs.

ALCF Resources

During CY 2018, the ALCF operated two INCITE production resources and two ALCC production resources, Mira and Theta. Mira is a 48K-node, 768K-core, 10-petaflops (PF) Blue Gene/Q with 768 TB of RAM. Mira mounts three general parallel file systems (GPFSs) with approximately 26.5 petabytes (PB) of usable space. Theta is a 4,392-node, 281K-core, 11.69-PF Cray XC40 with 891 TB of RAM. (The ALCF retains 2 racks of Theta as contingency to make up for failed parts elsewhere in the system.) Theta mounts one GPFS file system and one Lustre file system with approximately 8.6 PB and 9.2 PB of usable space, respectively. The filesystems on both Mira and Theta are also mounted by the visualization and analysis cluster Cooley. Both Mira and Theta have access to the facility-wide HPSS (high-performance storage system) tape archive.

The ALCF also operated two test and development Blue Gene/Q systems, Cetus and Vesta, and one test and development Cray XC40, Iota, in 2018. Cetus is a 4K-node, 64K-core Blue Gene/Q with 64 TB of RAM. Cetus shares file systems with Mira. Vesta is a 2K-node, 32K-core Blue Gene/Q with 32 TB of RAM. Vesta is an independent test and development resource and shares no resources with Mira or Cetus. Iota is a 44-node, 2,816-core Cray XC40 with 9 TB of RAM. Like Vesta, Iota is an independent test and development resource and shares no resources with Theta.

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

2.1 Resource Availability

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

This section reports on measures that are indicative of the stability of the systems and the quality of the maintenance procedures. In 2017, the ALCF, OLCF, and NERSC agreed that any event occurring during a scheduled maintenance that delays the return of a system to production by more than 4 hours will be counted as an adjacent unscheduled outage, as unscheduled availability, and as an additional interrupt.

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, the ALCF has agreed to metrics of 90 percent overall availability and 90 percent scheduled availability (in response to ASCR's request that all user facilities use a target of 90 percent for scheduled availability for the lifetime of the production resources). Table 2.2 summarizes the availability results.

Table 2.2 Availability Results for Theta

Theta (Cray XC40): 4,008-node, 257K-core, 64 TB MCDRAM, 770 TB DDR4				
	CY 2017		CY 2018	
	Target (%)	Actual (%)	Target (%)	Actual (%)
Scheduled Availability	90.0	98.6	90.0	98.0
Overall Availability	80.0	90.0	90.0	94.3

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 of Theta for the period January 1, 2018, through December 31, 2018, as annotated in Figure 2.1.

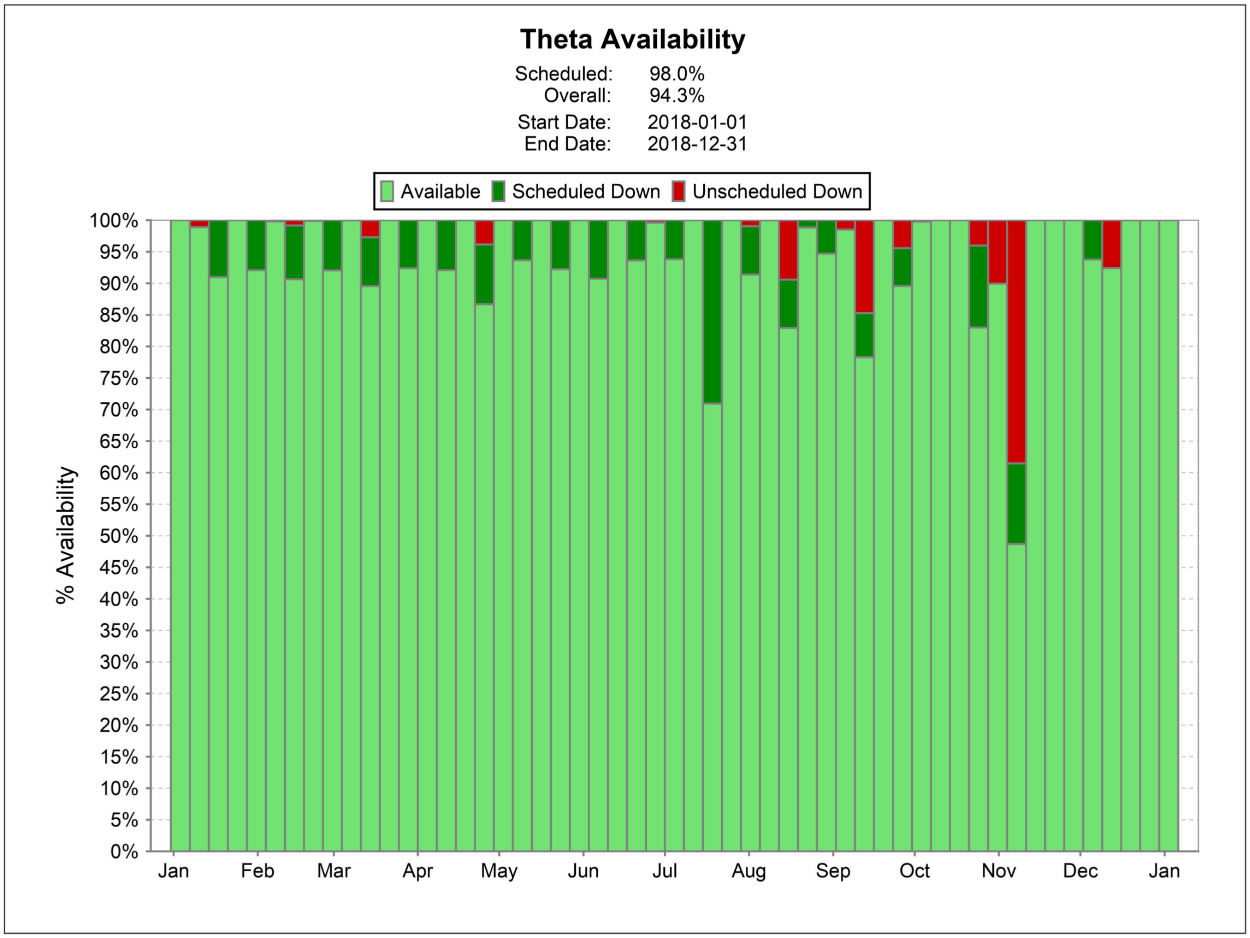

Figure 2.1 Theta Weekly Availability for CY 2018

Graph Description: Each bar in Figure 2.1 represents the average of seven days of core-hour usage. Each bar accounts for all of the time in one of three categories. The pale-green portion represents available core-hours; the darker green represents scheduled downtime for that week; and red represents unscheduled downtime. Dates are associated with the events and should be identifiable when cross-referencing with the graph. These unscheduled events are described in detail below.

March 12, 2018: DVS projected file system problem

An issue was discovered with the file system that the Cray Data Virtualization Service (DVS) was projecting, which caused 4,098 nodes to go down. A bug was filed with Cray, but it was not possible to gather enough information to determine a root cause. A new procedure is in place so that we can acquire more information from a single DVS server if it happens again. However, the issue has not occurred since, and with a major upgrade to the OS that was implemented mid-year, we do not expect this particular problem to reoccur.

April 23, 2018: Provisioning failed on logins

An issue arose a few hours after Theta was released from preventative maintenance work on April 23. During the maintenance, a command was run to copy over the Programming Environment (PE) image as a squashfs file system, instead of to a chroot file system. The login environment does not support this usage, which was not detected until Ansible (the configuration management system) failed after we were back in production. In an attempt to remedy the situation, the PE image was pushed out via the chroot method. Given the size of the PE image, the sync operations took several hours. Unfortunately, after all the sync operations finished, Ansible failed again. Although some time was spent in an attempt to diagnose the issue, a root cause could not be determined. At approximately 01:30 a.m., the team decided that the various diagnostic efforts had put the system in a non-salvageable state and that the best path forward would be to redeploy the images from scratch. As this is a long process and the team was exhausted, further actions were deferred until the following business day. The next business day, the team was able to redeploy the images from scratch and resume normal processes. Improvements have been made to the process to prevent this from happening again.

August 12, 2018: Scheduling issue between ALPS and Cobalt
On August 12, the Cray Application Level Placement Scheduler (ALPS) started sending unexpected messages to Cobalt, which caused Cobalt to fail in updating nodes as free. This event occurred because the cleanup detection in Cobalt failed due to messages it could not parse. As a result, resources (nodes) would not free up when jobs ended; thus, we ran out of nodes that were marked as able to run jobs. The unexpected messages issue cleared up after nearly 16 hours. More debugging features were added to Cobalt so that there is a much greater chance of debugging this issue in the future if it were to happen again (which has not occurred).

September 11–12, 2018: Substation breaker tripped
At approximately 3:20 p.m. on September 11, 2018, a breaker in the substation feeding Theta tripped, thereby removing power from one of Theta's wall panels. Six Theta compute cabinets in row 1 and one blower cabinet lost power. The loss of the blower cabinet caused the remainder of row 1 to be sent an emergency power-off (EPO) signal. The loss of all of row 1 triggered an Aries interconnect routing recalculation, which unexpectedly failed, resulting in the loss of compute communications for the remaining row (row 0). These events effectively rendered Theta unusable. After an initial investigation into the power loss with no obvious causes, facilities management decided to completely power down the substation to do further testing overnight. As such, Theta was powered down for the night. After testing was concluded, Theta was brought back into service at approximately 4:00 p.m. on September 12, 2018. Upon further investigation, facilities management determined that it is a "high probability" that the circuit breaker and trip unit were the cause for the power outage. As a result, this breaker would be replaced during a future maintenance window when a replacement breaker was available.

September 29, 2018: Substation breaker tripped again
US-B-6 breaker tripped again (see the September 11–12 event above) before a replacement breaker was obtained. It tripped in exactly the same manner, which confirmed it as a suspect trip unit. No replacement was yet available. However, the facility manager confirmed taking steps to expedite acquiring a replacement. The facility manager also suggested an interim fix, but it would require a subsequent outage. We decided to put Theta back in production as there was a chance it would survive until the replacement breaker arrived; however, if it failed again, then we would execute the interim fix, which would be ready to deploy as soon as an outage started. Power was restored at 3:30 p.m. — an approximately 9-hour outage. Fortunately, the breaker worked fine until a planned preventative maintenance to install a new but temporary breaker.

This temporary breaker was an off-the-shelf component that was rated for about 80% of the original breaker but was more than enough to run the Theta equipment. After the preventative maintenance window on November 5 (see item), the final replacement breaker was installed. Facility management now has two spare breakers ― one identical to the current breakers and the off-the-shelf breaker that could be used under the right circumstances.

October 28, 2018: File system with Postgres database (DB) full
The file system that the Postgres database for the HPSS resides on ran out of space. Upon detection, work was performed to clean up space, including cleaning up unused records in the database. Theta was returned to normal service after space was cleaned up, and the database was verified to be working correctly.

November 5–8, 2018: Substation maintenance turns into multiday outage
On the regular preventative maintenance day of November 5, the new final replacement breaker for US-B-6 was installed. While US-B-6 was powered down, the facility manager performed a scheduled maintenance on the substation servicing the main breaker and switchboard and tested the health of the transformers. This testing of the high-voltage transformer revealed low resistivity based on generic standards for NETA (the international electrical testing association) standards; low resistivity could allow arcing in the transformer. The facility manager gathered an oil sample and sent it to a NETA testing facility for dissolved gas analysis to determine whether arcing had already occurred. Because of the safety concerns over arcing, the substation was kept powered down, an outage that the ALCF had not planned on taking. The results of the testing came back late November 7. Facility management had the substation back online on November 8, and Operations brought Theta back up immediately afterward, such that it was released to users by 5:00 p.m. on November 8. Given that the manufacturer's tolerances allowed more than the generic NETA standards, the substation was thus determined to be operating normally.

December 14, 2018: Water line break
The Argonne Chilled Water Plant ceased operations for nearly 4 hours to repair a broken water supply line first detected around 12:30 a.m. The plant resumed supplying chilled water at 4:15 a.m.; however, the rising temperatures in the data center caused many ALCF resources to automatically turn themselves off, including Theta (and Mira). The entire ALCF infrastructure and both supercomputers had to be restarted. Theta was fully restored and released to users at 9:17 p.m. All Theta hardware (compute and storage) survived this event.

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 the current MTTI and MTTF values.

Table 2.3 MTTI and MTTF Results for Theta

Theta (Cray XC40): 4,008-node, 257K-core, 64 TB MCDRAM, 770 TB DDR4				
	CY 2017		CY 2018	
	Target	Actual	Target	Actual
System MTTI	N/A [a]	6.94 days	N/A	10.12 days
System MTTF	N/A	30.30 days	N/A	32.55 days

[a] N/A = Not applicable.

Theta currently functions on a biweekly maintenance schedule. The ALCF takes the machine out of service every other Monday to perform Cray driver upgrades, hardware replacements, OS upgrades, etc. While Theta is out of service concurrently with Mira, the ALCF uses that opportunity to perform other potentially disruptive maintenance such as facilities power and cooling work, and storage systems upgrades and patching. Thus, Theta's biweekly maintenance schedule caps MTTI at 14 days but does not directly affect MTTF.

2.2 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.

The following sections discuss system allocation and usage, total system utilization percentage, and capability usage for Theta. 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).

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 utilization results, and Figure 2.2 shows system utilization over time by program.

Table 2.4 System Utilization Results for Theta

Theta (Cray XC40): 4,008-node, 257K-core, 64 TB MCDRAM, 770 TB DDR4				
	CY 2017		CY 2018	
	Target	Actual	Target	Actual
System Utilization	N/A [a]	82.7%	N/A	92.1%

[a] N/A = Not applicable.

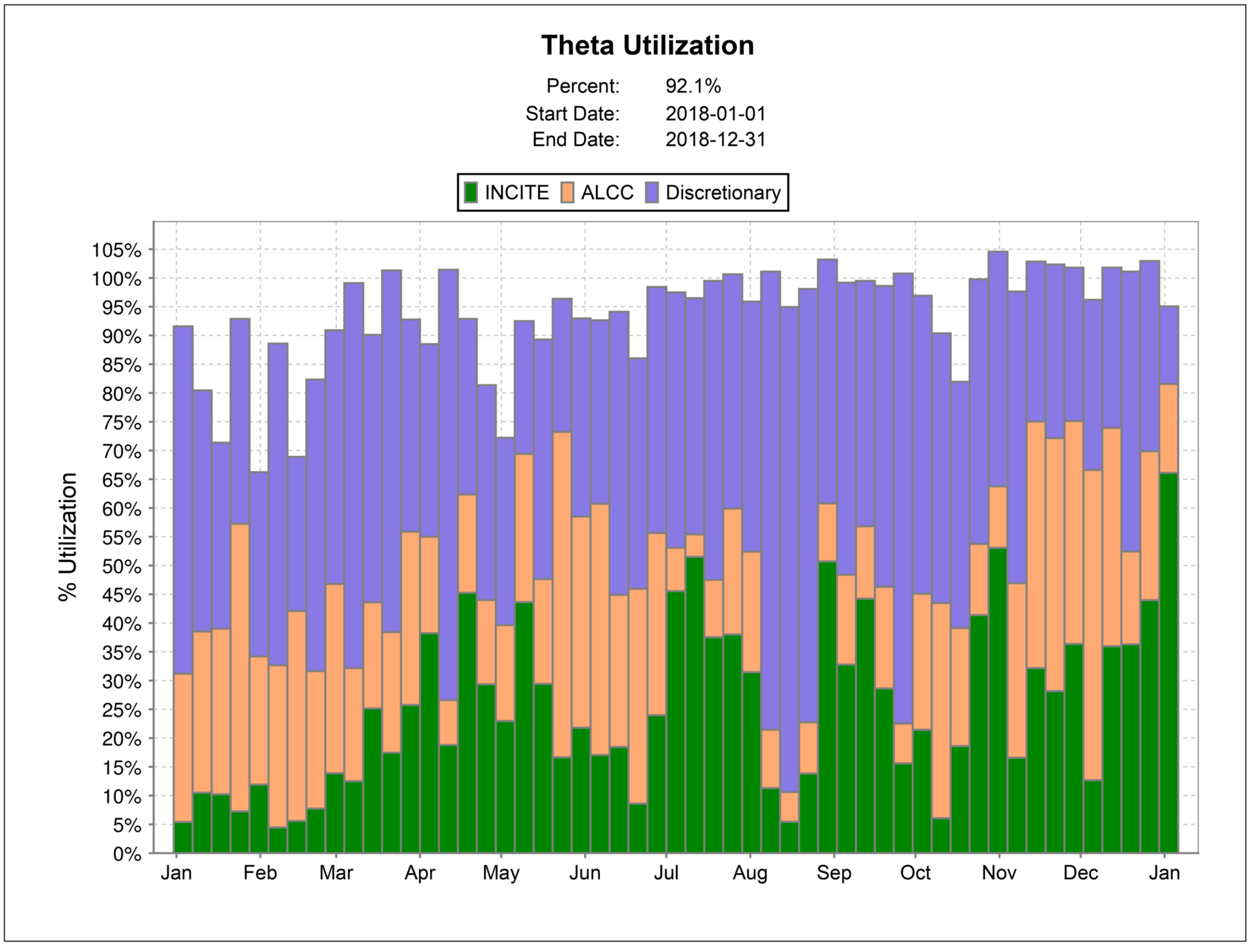

Figure 2.2 System Utilization over Time by Program on Theta

The system utilization for Theta was 92.1 percent for its 2018 production period of January 1, 2018, through December 31, 2018.

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 the ALCF's DOE Program Manager determines the hours available. Of the hours available in 2018, the first half of the year was allocated based on 30 percent to the INCITE program, 30 percent to the ALCC program allocations, and 40 percent to DD allocations; and for the second half of the year, the rates were 30 percent to the INCITE program, 20 percent to the ALCC program allocations, and 50 percent to the 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 2018, Theta delivered a total of 2.0 billion core-hours.

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

Theta (Cray XC40): 4,008-node, 257K-core, 64 TB MCDRAM, 770 TB DDR4						
	CY 2017			CY 2018		
	Allocated	Used		Allocated [b]	Used	
	Core-hours	Core-hours	% of total used	Core-hours	Core-hours	% of total used
INCITE	N/A [a]	N/A	N/A	390.9M	529.4M	27.1
ALCC	209.5M	84.9M	11.7	372.0M	508.2M	26.0
DD	628.2M	222.4M	30.7	1.5B	916.6M	46.9
Total	1.4B	724.1M	100.0	2.3B	2.0B	100.0

[a] N/A = Not applicable.

[a] Theta allocation amounts for 2018 INCITE program and Jan. 1-June 30, 2018 of the 2017/2018 ALCC program were based on a smaller machine size due to the late-2017 timing of the machine expansion.

Summary: For CY 2018, 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. Table 2.6 and Figure 2.3 show that the ALCF has substantially exceeded these metrics set for INCITE. Although no targets are set, data are also provided in the table for ALCC and DD projects as reference, and Figure 2.4 shows the overall distribution of job sizes over time.

Table 2.6 Capability Results for Theta

Theta (Cray XC40): 4,008-node, 257K-core, 64 TB MCDRAM, 770 TB DDR4						
	CY 2017			CY 2018		
Capability Usage	Total Hours	Capability Hours	Percent Capability	Total Hours	Capability Hours	Percent Capability
INCITE Overall	N/A [a]	N/A	N/A	529.4M	340.8M	64.4
INCITE High	N/A	N/A	N/A	529.4M	70.2M	13.3
ALCC Overall	84.9M	37.6M	44.3	508.2M	139.8M	27.5
ALCC High	84.9M	1.0M	1.2	508.2M	18.3M	3.6
Director's Discretionary Overall	222.4M	69.1M	31.1	916.6M	337.4M	36.8
Director's Discretionary High	222.4M	26.4M	11.9	916.6M	98.9M	10.8
TOTAL Overall	724.1M	290.7M	40.1	2.0B	818.0M	41.9
TOTAL High	724.1M	135.1M	18.7	2.0B	187.4M	9.6

[a] N/A = Not applicable.

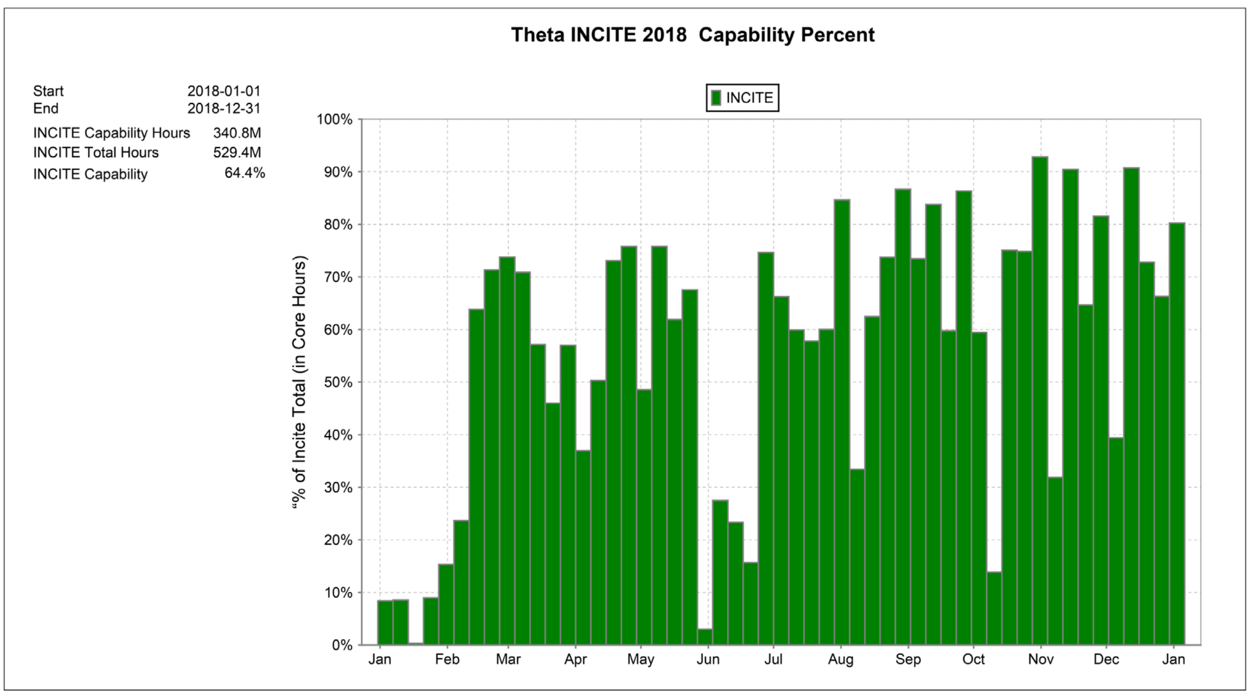

Figure 2.3 Theta Overall Capability for INCITE

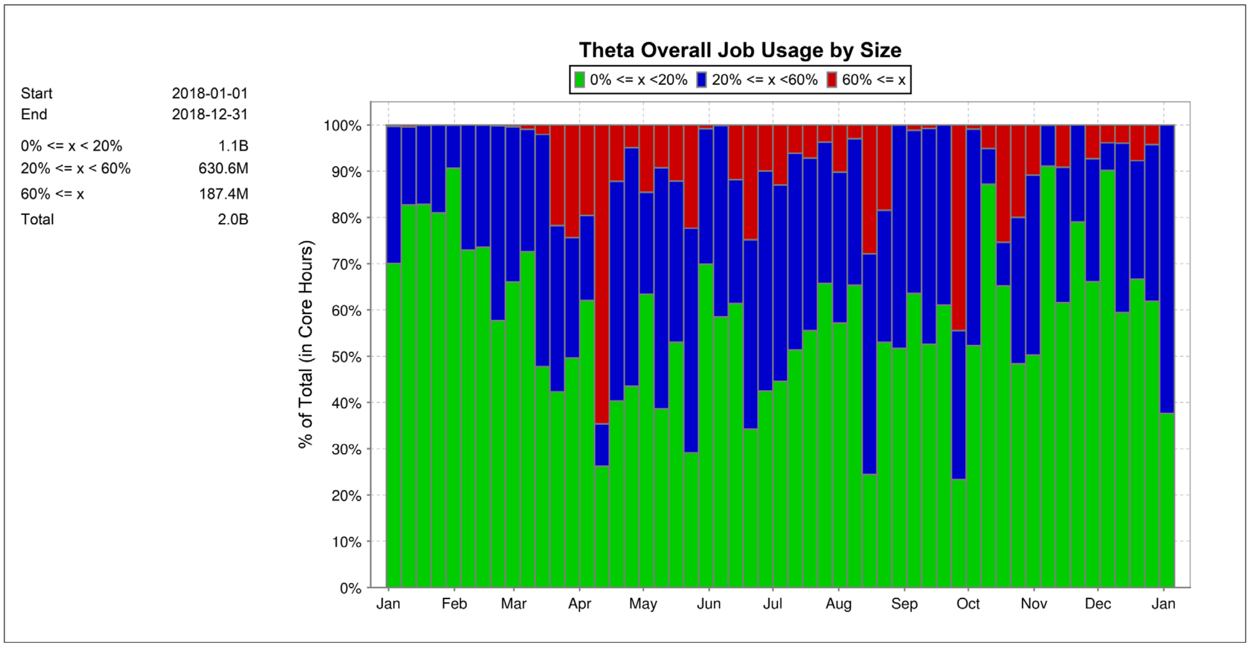

Figure 2.4 Theta Overall Job Usage by Size

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, the ALCF has agreed to metrics of 90 percent overall availability and a target of 90 percent scheduled availability (in response to ASCR's request that all user facilities use a target of 90 percent for scheduled availability for the lifetime of the production resources). Table 2.7 summarizes the availability results.

Table 2.7 Availability Results for Mira

Mira (Blue Gene/Q): 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	CY 2017		CY 2018	
	Target (%)	Actual (%)	Target (%)	Actual (%)
Scheduled Availability	90.0	99.5	90.0	99.3
Overall Availability	90.0	96.4	90.0	96.5

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 of Mira for the period January 1, 2018, through December 31, 2018, as annotated in Figure 2.5.

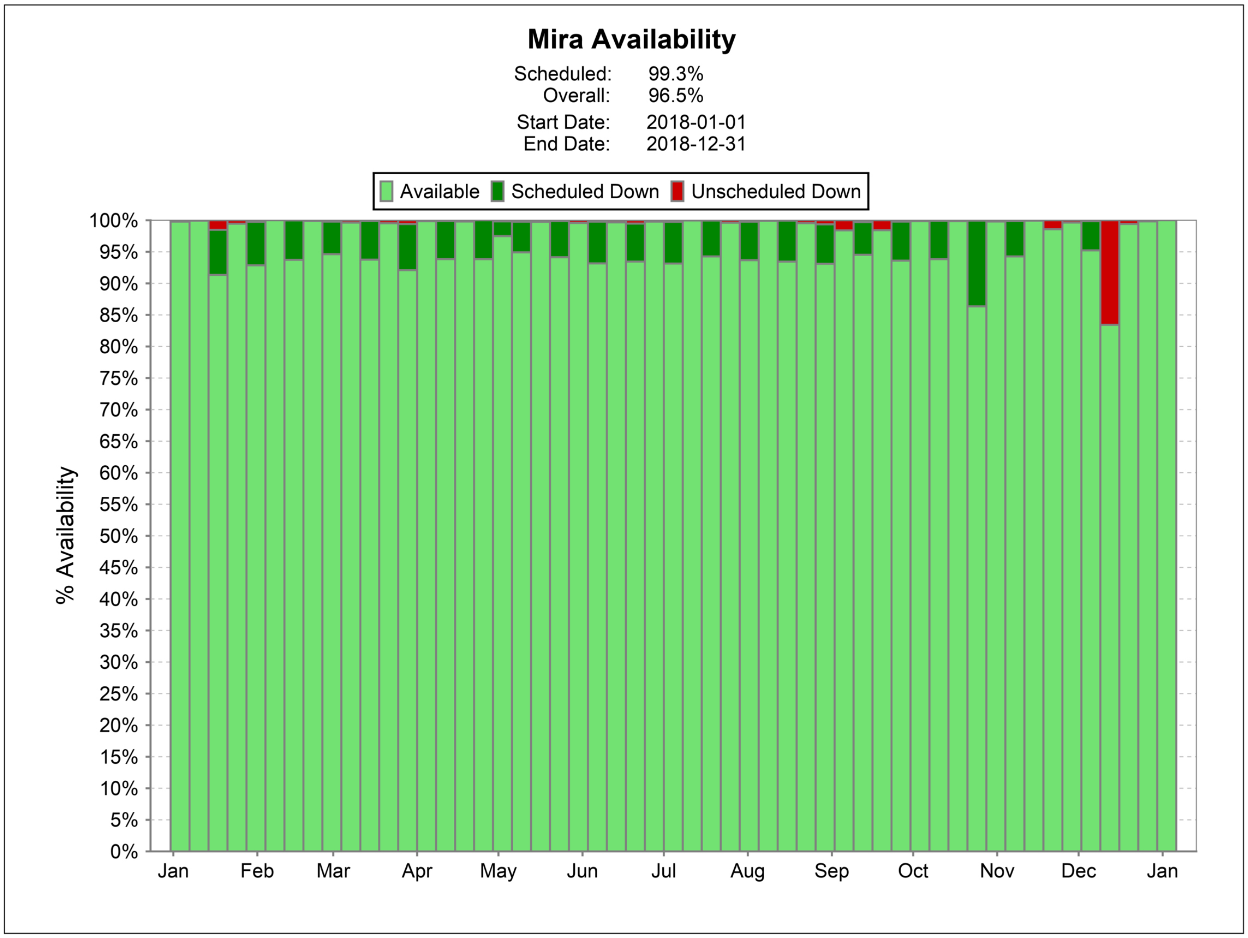

Figure 2.5 Mira Weekly Availability for CY 2018

Graph Description: 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. Dates are associated with the events and should be identifiable when cross-referencing with the graph. These unscheduled events are described in detail below.

January 15, 2018: IONs rebooted late in preventative maintenance period

Mira required the IONs to be rebooted late during the preventative maintenance period causing a delayed return to service.

September 7, 2018: runjob_server threads segfaulting

The Mira runjob_server worker threads were segfaulting, resulting in a lack of responsiveness and an inability to boot blocks and start jobs. This incident was traced to the DB2 database, which was allocating an inordinate amount of memory on the service node. A couple of fixes were implemented: (1) the DB2's memory allocation was reduced, and (2) the runjob_server was changed to limit the number of worker threads.

September 19–20, 2018: GPFS waiters
On September 19, the mira-fs1 file system was experiencing long delays in response to user commands. This situation is usually a result of "long-waiters" in GPFS (I/O commands waiting longer than normal because of extreme amounts of contention). The administrators held any jobs that would use mira-fs1, but let jobs using mira-fs0 continue to run. However, the problems seemed to be affecting not only GPFS cluster nodes but also the Mira ION nodes. The administrators decided to drain Mira with a reservation. Once all of the jobs had drained, the long waiters cleared up. Because the problems cleared up, the network shared disk (NSD) nodes were not rebooted, although the administrators rebooted the Mira ION nodes. The mira-fs1 long-waiters situation occurred because the mira-fs1 file system was low on free metadata blocks. Although the reservation cleaned up the system, the underlying problem was a scarcity of metadata blocks. A fix to the problem, implemented after this event, was the addition of metadata logical unit numbers (LUNs) to the mira-fs1 system.

November 21, 2018: /tmp and /dev/shmem full
A user running the NWChem application filled both /tmp and /dev/shmem at the same time, bringing down Mira mid-planes. Scheduling was halted for 2 hours prior to the Thanksgiving holiday, and as a result, the administrator has placed greater restrictions on the IONs and mount-points so that the event will not occur again. With this change, the mount-points can still be used, but with limitations on how much memory a user can use in total on the IONs.

December 14, 2018: Water line break
A chilled water supply line broke that forced the Argonne Chilled Water Plant to shut down, given that the water lost could not be resupplied fast enough by the make-up water. The break was detected around 12:30 a.m. and was fixed by Argonne around 4:15 a.m. The Chilled Water Plant was providing chilled water by 4:15 a.m. Many of the ALCF resources shut themselves down due to high temperatures in the data center, as did Mira (and Theta). The entire ALCF infrastructure had to be restarted, as well as the supercomputers. Mira was mostly restored and returned to users at 4:46 a.m. the following day. There were a few hardware components that still required service when it was returned — the result being that some of the largest partitions were unavailable until all of the repairs were complete, which was later that day.

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.

Table 2.8 MTTI and MTTF Results for Mira

Mira (Blue Gene/Q): 48K-node, 768K-core, 1.6 GHz, 768 TB RAM				
	CY 2017		CY 2018	
	Target	Actual	Target	Actual
System MTTI	N/A [a]	10.37 days	N/A	12.18 days
System MTTF	N/A	36.44 days	N/A	72.73 days

[a] N/A = Not applicable.

Mira currently functions on a biweekly maintenance schedule. The 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, the ALCF uses that opportunity to perform other potentially disruptive maintenance such as facilities power and cooling work, and storage systems upgrades and patching. Thus, the 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 Mira. 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 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 2017		CY 2018	
	Target	Actual	Target	Actual
System Utilization	N/A [a]	95.5%	N/A	93.6%

[a] N/A = Not applicable.

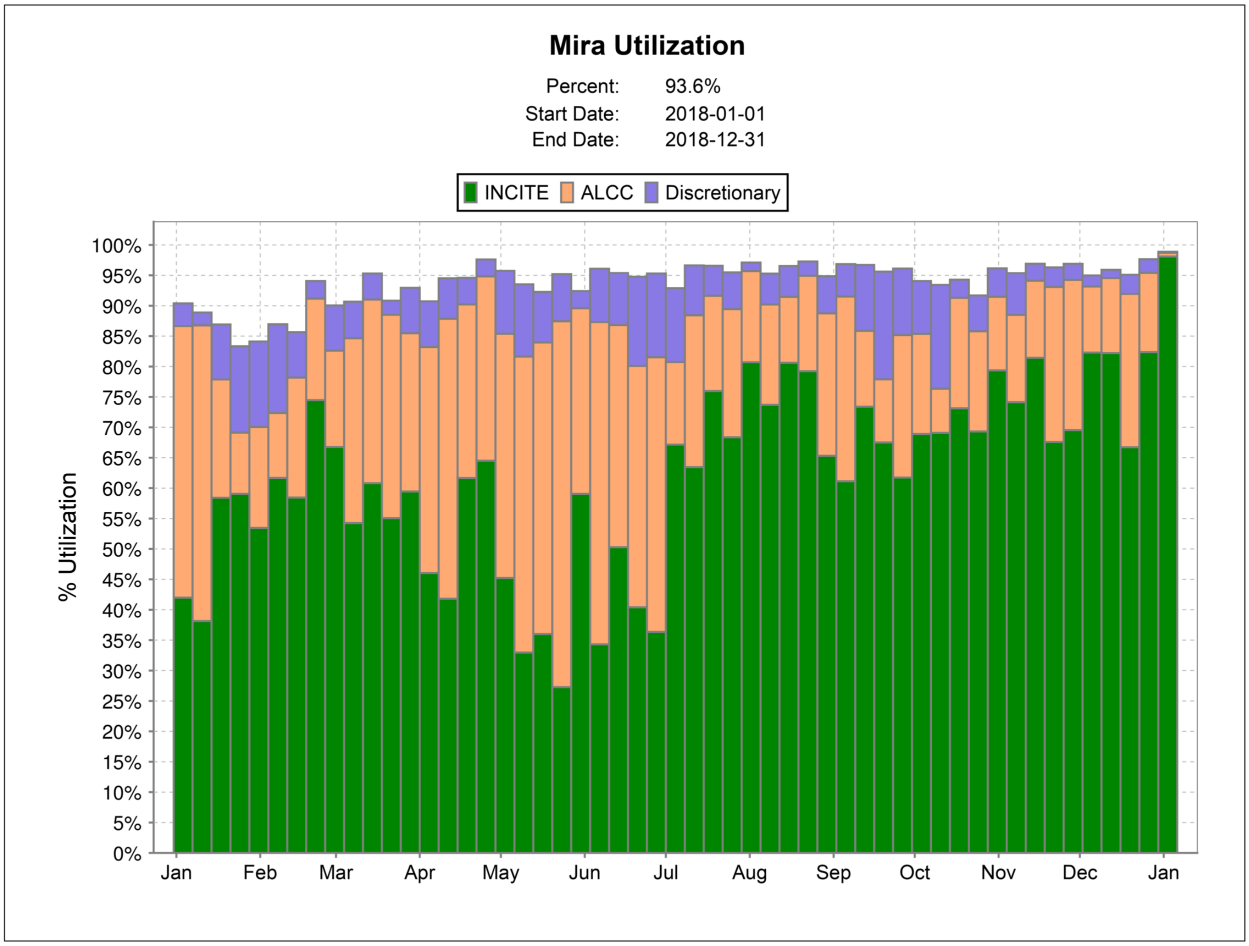

Figure 2.6 Mira System Utilization over Time by Program

The system utilization for Mira was 93.6% for its 2018 production period of January 1, 2018, through December 31, 2018.

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 the ALCF's DOE Program Manager determines the hours available. Of the hours available, 60 percent are allocated to the INCITE program, up to 20 percent are available for the ALCC program allocations, and 20 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 2018, Mira delivered a total of 6.3 billion core-hours.

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 2017			CY 2018		
	Allocated	Used		Allocated	Used	
	Core-hours	Core-hours	%	Core-hours	Core-hours	%
INCITE	3.5B	4.5B [a]	70.3	3.4B	4.1B [e]	65.8
ALCC	1.8B	1.6B [b]	24.9	1.5B	1.7B [f]	26.6
DD	1.7B	309.6M [c]	4.9	1.4B	472.2M [g]	7.6
Total	7.0B	6.4B [d]	100	6.2B	6.3B [h]	100.0

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

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

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

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

[e] Usage includes 0.3M core-hours from Cetus production jobs.

[f] Usage includes 0.5M core-hours from Cetus production jobs.

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

[h] Usage includes 11.1M core-hours from Cetus production jobs.

Summary: For CY 2018, 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 racks or more (16.7 percent), and High Capability requires a minimum of 10 percent of the INCITE core-hours be run on 16 racks or more (33.3 percent). See Appendix A for more detail on the capability calculation. Table 2.11 and Figure 2.7 show that the 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.

Table 2.11 Capability Results for Mira

Mira (Blue Gene/Q): 48K-node, 768K-core, 1.6 GHz, 768 TB RAM						
	CY 2017			CY 2018		
Capability Usage	**Total Hours**	**Capability Hours**	**Percent Capability**	**Total Hours**	**Capability Hours**	**Percent Capability**
INCITE Overall	4.5B [a]	3.2B	72.3	4.1B [e]	2.9B	70.8
INCITE High	4.5B [a]	1.2B	27.5	4.1B [e]	904.4M	22.0
ALCC Overall	1.6B [b]	1.0B	65.7	1.7B [f]	888.6M	53.3
ALCC High	1.6B [b]	262.7M	16.6	1.7B [f]	261.4M	15.7
Director's Discretionary Overall	309.6M [c]	155.3M	50.2	472.2M [g]	254.7M	53.9
Director's Discretionary High	309.6M [c]	65.9M	21.3	472.2M [g]	66.6M	14.1
TOTAL Overall	6.4B [d]	4.4B	69.5	6.3B [h]	4.1B	64.9
TOTAL High	6.4B [d]	1.6B	24.5	6.3B [h]	1.2B	19.7

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

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

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

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

[e] Usage includes 0.3M core-hours from Cetus production jobs.

[f] Usage includes 0.5M core-hours from Cetus production jobs.

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

[h] Usage includes 11.1M core-hours from Cetus production jobs.

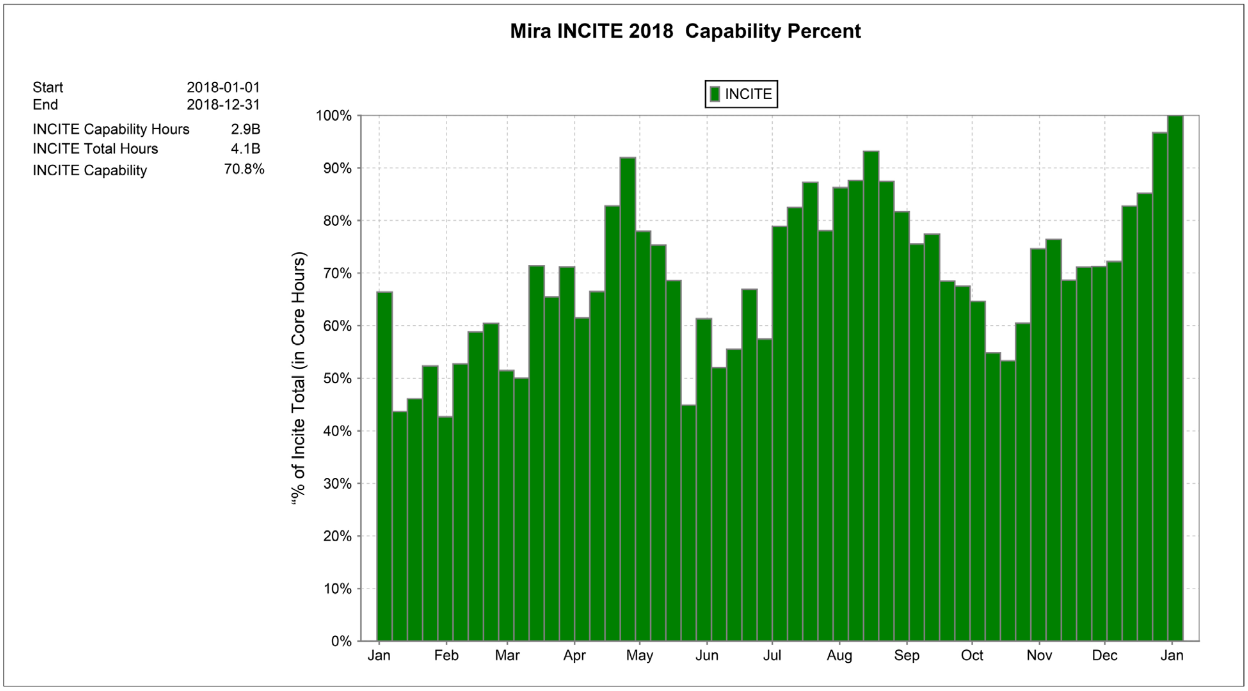

Figure 2.7 Mira Overall Capability for INCITE

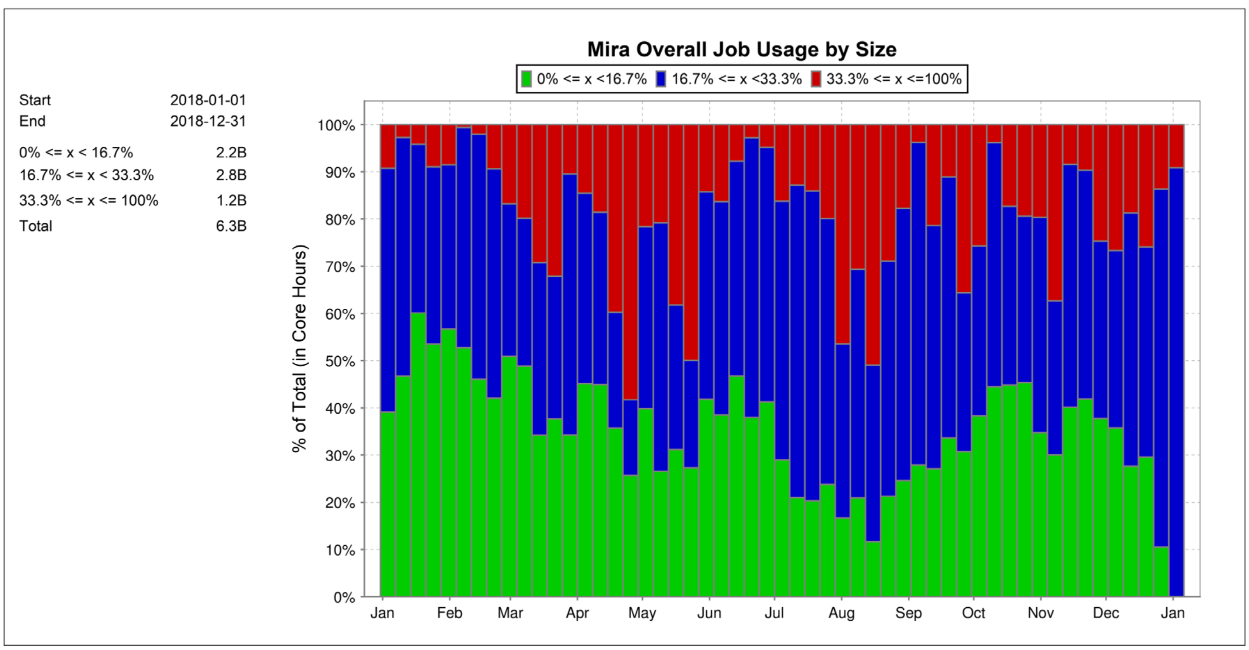

Figure 2.8 Mira Overall Job Usage by Size

Conclusion

The ALCF is maximizing the use of its HPC systems and other resources consistent with its mission. We have exceeded our targets for the metrics of system availability, INCITE hours delivered, and capability hours delivered. The ALCF closely tracks hardware and software failures and their impact(s) on user jobs and metrics. These data are used as a factor in the selection of troubleshooting efforts and improvement projects. In CY 2018, as in previous years, this regular failure analysis has continued to drive code improvements to Cobalt, the ALCF's job scheduler; to our configuration management systems; and to the debugging of storage system problems.

This page intentionally left blank.

Section 3. Allocation of Resources

Is the allocation of resources reasonable and effective?

ALCF Response

As our results show in Section 8, these are a reasonable allocation of resources. Here are a few data points the ALCF looks at when analyzing usage statistics for the various allocation programs.

3.1 Allocation of Resources

3.1.1 Usage of the INCITE and ALCC Hours

The INCITE 2018 program allocated 3.4 billion core-hours on Mira and 390.9 million core-hours on Theta. The allocation usage levels are shown for individual projects in Figure 3.1 for Mira and in Figure 3.2 for Theta. Of the 32 INCITE projects sharing the two major resources, 26 projects used more than 90 percent of their allocation. Twenty projects used their entire allocation (or more), including 5 projects using more than 150 percent, of which 2 projects used more than 200 percent. These projects used the extra core-hours to achieve additional milestones. The overuse of Mira and Theta 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.

Of the remaining 12 projects with outstanding allocations, 9 projects used more than 70 percent of their time. Only 3 projects used less than 50 percent of their allocations. A total of 4.1 billion core-hours was delivered to INCITE on Mira (Table 3.1). The contribution from Cetus to the total INCITE hours delivered was 300,000 core-hours.

For Theta's 14 INCITE projects, 8 projects used more than 90 percent of their allocations. Seven projects used their entire allocations (or more), including 2 projects using more than 200 percent, leading to 529.4M core-hours used on Theta in 2018 (Table 3.2).

For the 2017–2018 ALCC year, 22 projects had allocations on Mira for a total of 1.7 billion core-hours and 10 projects on Theta for a total of 419 million core-hours. The allocation usage levels are shown in Figure 3.3 for Mira and in Figure 3.5 for Theta. Eighteen 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 their usage and less than 0.1 percent of the total ALCC 2017–2018 usage. Of the remaining 9 projects, 1 used 75 percent of its allocation or more and only 5 projects used less than 50 percent of their allocations.

The 2018–2019 ALCC year is approximately halfway through its allocation cycle. So far, 21 unique projects have received allocations of 1.9 billion core-hours across Mira and Theta. The projects have used a total of 780.8 million core-hours from July 1, 2018 through December 31, 2018. The allocation usage levels for Mira and Theta are shown in Figures 3.4 and 3.6, respectively. Table 3.3 shows ALCC hours allocated and used on Mira in CY 2018. Table 3.4 shows ALCC hours allocated and used on Theta in CY 2018.

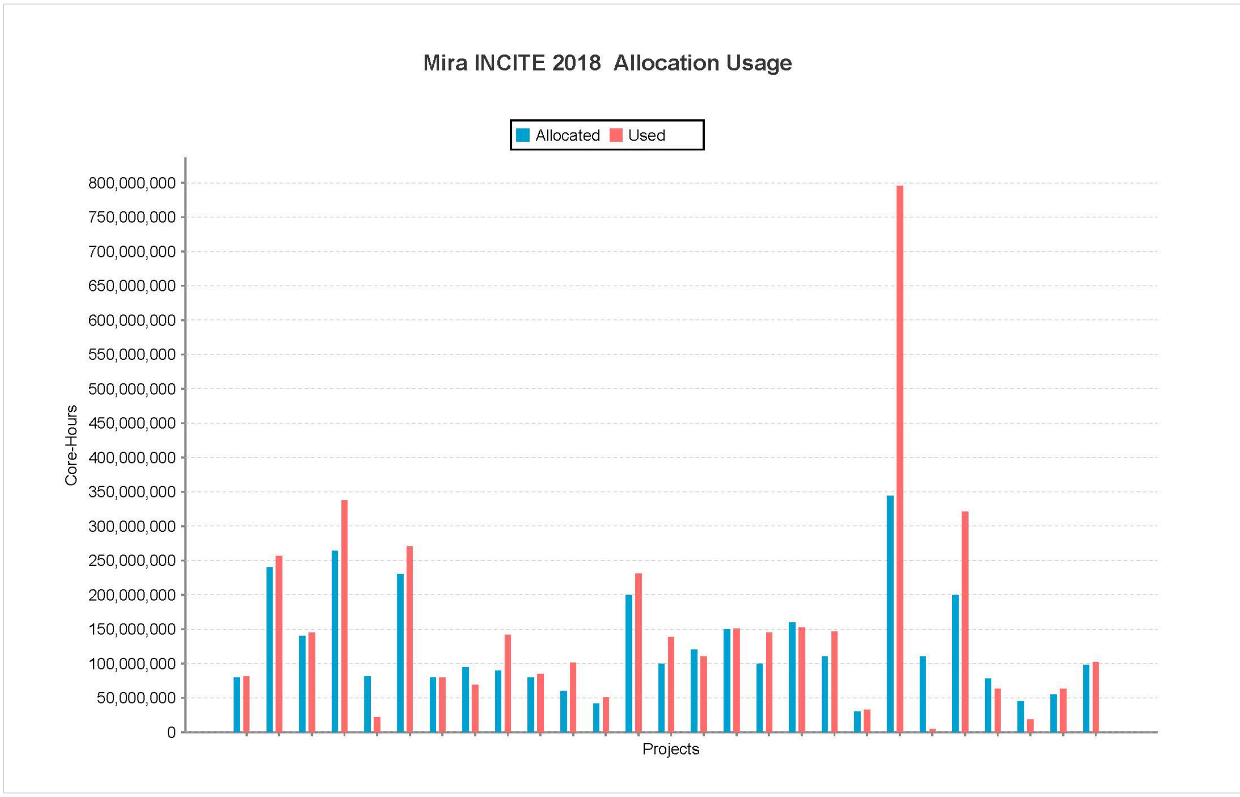

Figure 3.1 Mira INCITE 2018 Allocation Usage *

* Projects are randomly ordered.

Table 3.1 INCITE 2018 Time Allocated and Used on Mira

Projects	Mira
Allocated Core-Hours	3.4B
Used Core-Hours	4.1B [a]

[a] Usage includes 0.3M core-hours from Cetus production jobs.

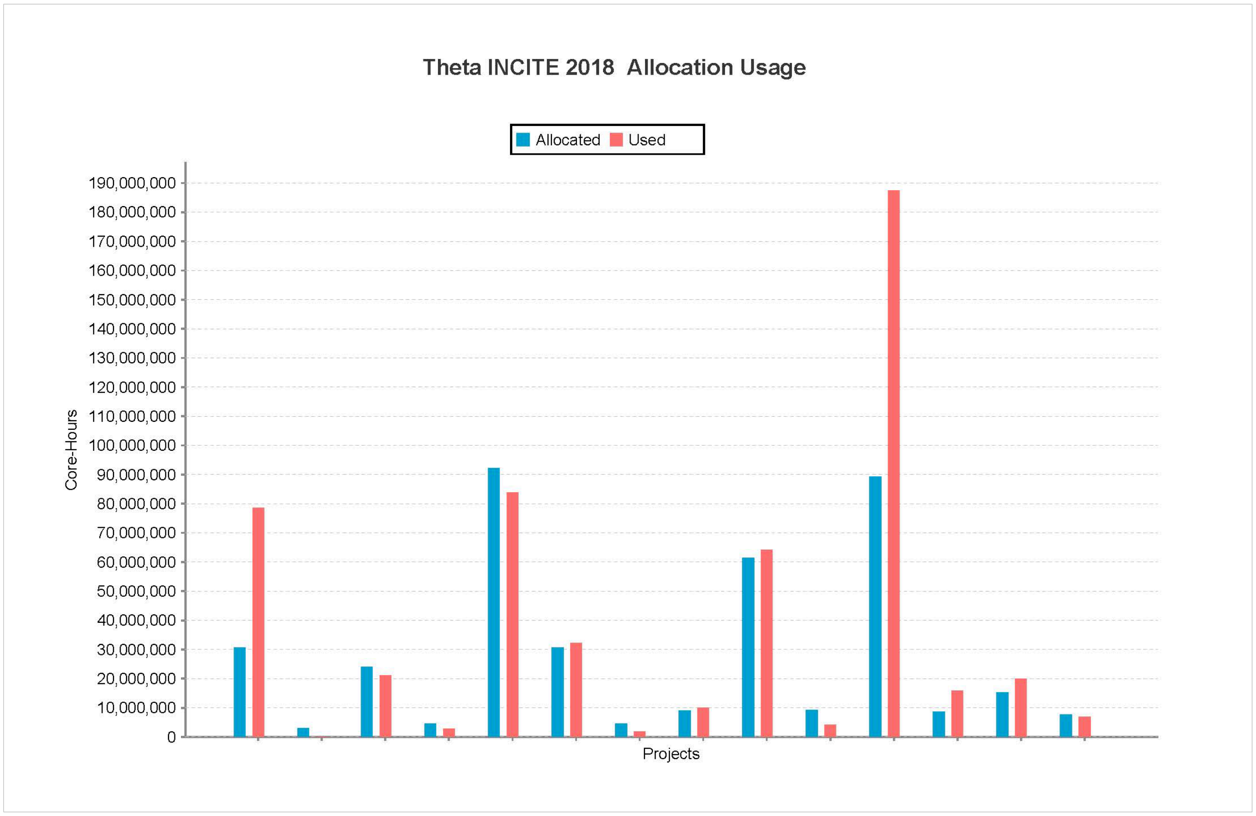

Figure 3.2 Theta INCITE 2018 Allocation Usage *

* Projects are randomly ordered.

Table 3.2 INCITE 2018 Time Allocated and Used on Theta

Projects	Theta
Allocated Core-Hours	390.9M
Used Core-Hours	529.4M

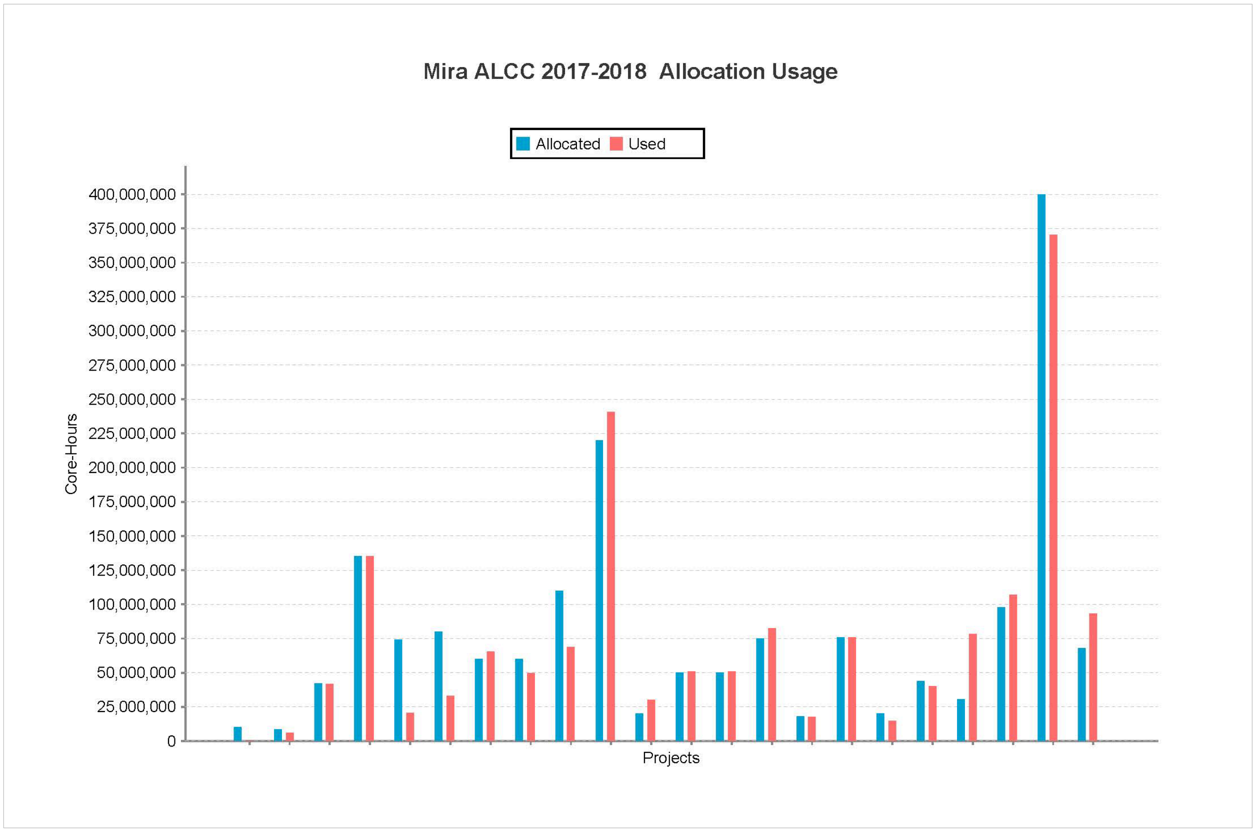

Figure 3.3 Mira ALCC 2017–2018 Allocation Usage *

* Projects are randomly ordered.

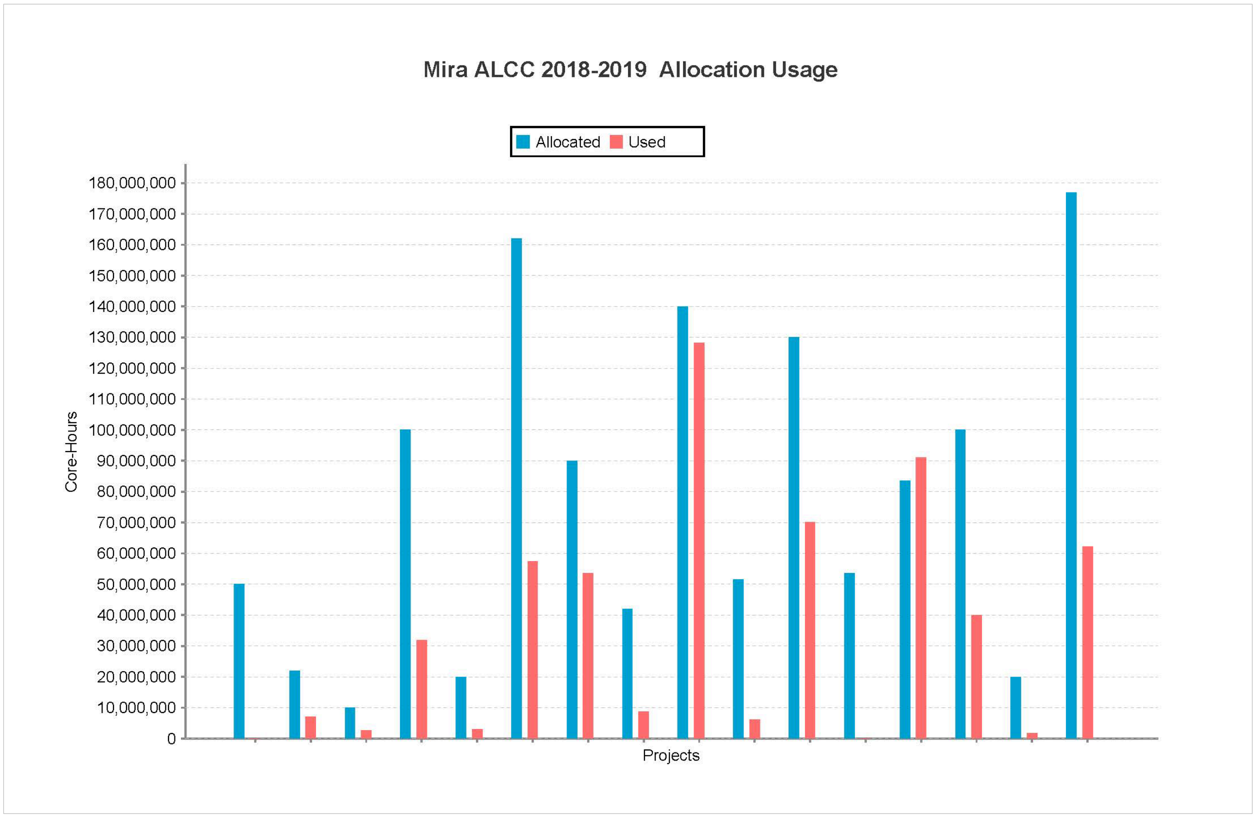

Figure 3.4 Mira ALCC 2018–2019 Allocation Usage *

* Projects are randomly ordered.

Table 3.3 ALCC Time Allocated and Used on Mira in CY 2018

Projects	Mira
Allocated Core-Hours	1.5B
Used Core-Hours	1.7B [a]

[a] Usage includes 0.5M core-hours from Cetus production jobs.

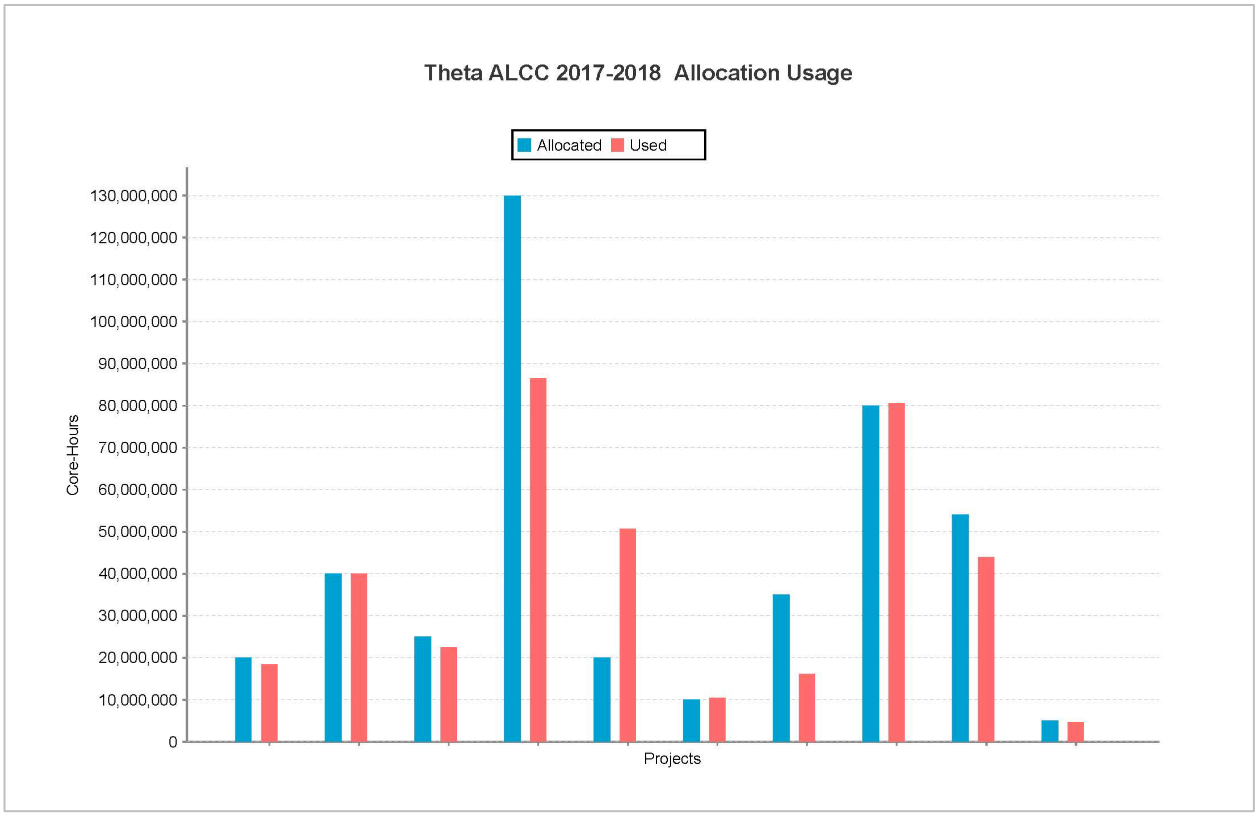

Figure 3.5 Theta ALCC 2017–2018 Allocation Usage *

* Projects are randomly ordered.

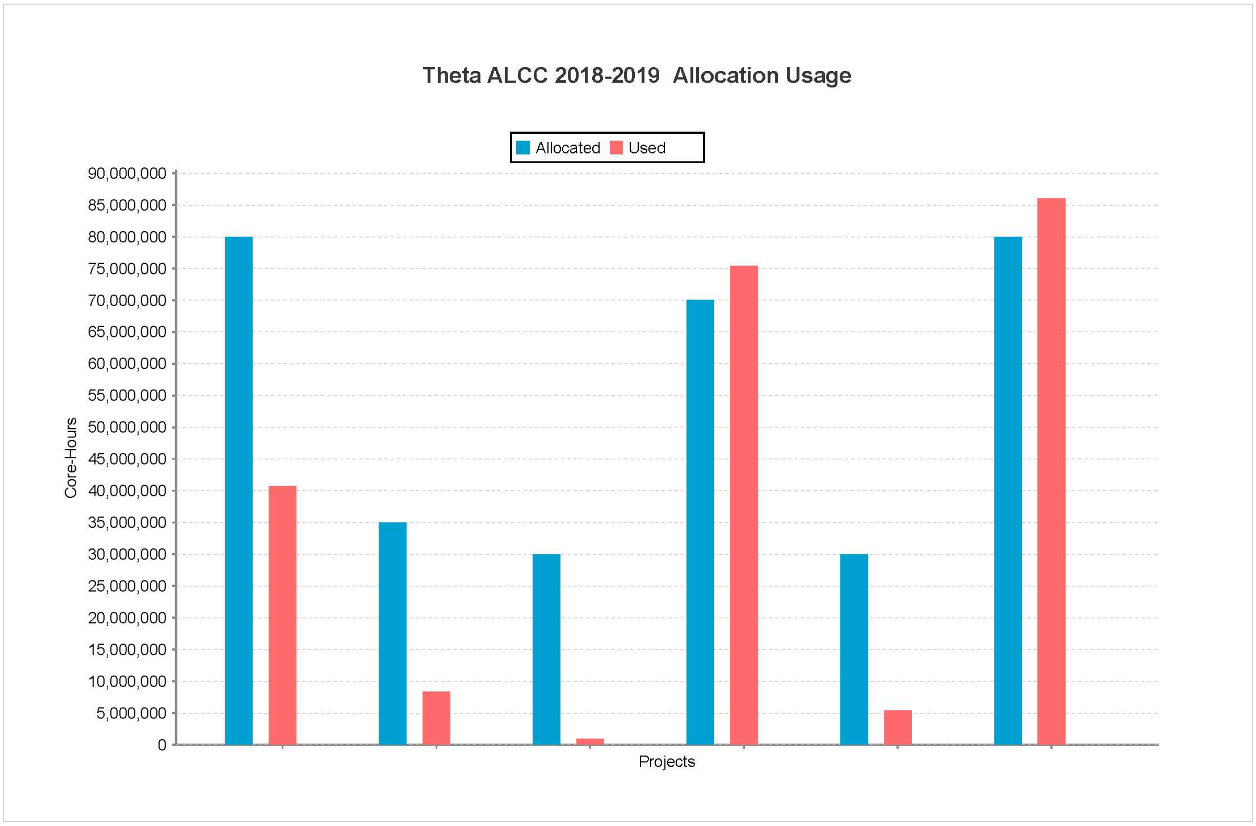

Figure 3.6 Theta ALCC 2018–2019 Allocation Usage *

* Projects are randomly ordered.

Table 3.4 ALCC Time Allocated and Used on Theta in CY 2018

Projects	Theta
Allocated Core-Hours	372.0M
Used Core-Hours	508.2M

3.1.2 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 five categories:

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

INCITE and ALCC proposal preparation allocations are offered for projects that are targeting submission of an INCITE or ALCC 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 the ALCF.

The 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 and the Data Science Program.

Internal/support projects are devoted to supporting the ALCF mission. The 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.

As of the midpoint of 2018, the discretionary pool grew to 20 percent of the system to support the needs of the ECP program. ECP and the computing facilities run a Resource Allocations Council (RAC) that meets on a monthly basis to discuss the computing needs for ECP. Needs allocations are discussed in the council and brought forward to the ALCF. The program targets approximately 10 percent of the total time, or half of the 20 percent discretionary pool.

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 amounts cannot be accommodated.

Table 3.5 and Table 3.6 show the total DD time allocated and used on Theta and Mira, respectively, during 2018. By its very nature, the DD program is amenable to over-allocation because often time is left unused; however, it should be noted that these totals do not represent open allocations for the entire calendar year. A project might have a 1-million core-hour allocation that only persists for three months, but that 1-million core-hour allocation is counted entirely in the annual total core-hour number. Projects are not guaranteed the allocated time; rather, the time is provided on a first-come, first-served basis. DD projects run at a lower priority than INCITE or ALCC projects, which reduces the amount of time available for their use. Exceptions are made for some internal projects that support acceptance of new hardware or support of users, which is in line with the ALCF core mission.

Table 3.5 DD Time Allocated and Used on Theta

Projects	Theta
Allocated Core-Hours	1.5B
Used Core-Hours	916.6M

Table 3.6 DD Time Allocated and Used on Mira

Projects	Mira
Allocated Core-Hours	1.4B
Used Core-Hours	472.2M [a]

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

Lists of the CY 2018 DD projects, including title, PI, institution, and hours allocated, are provided in Appendix B (for Mira) and Appendix C (for Theta).

Figures 3.7 and 3.8 provide a breakdown of the CY 2018 DD allocations by domain for Mira and Theta, respectively.

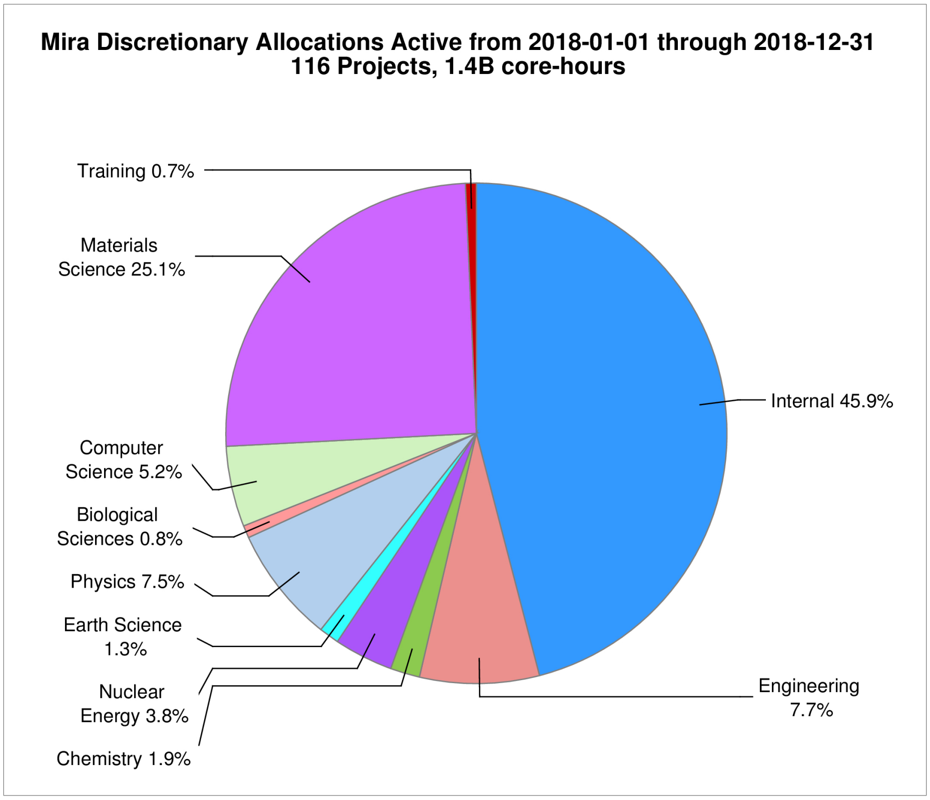

Figure 3.7 Mira CY 2018 DD Allocations by Domain

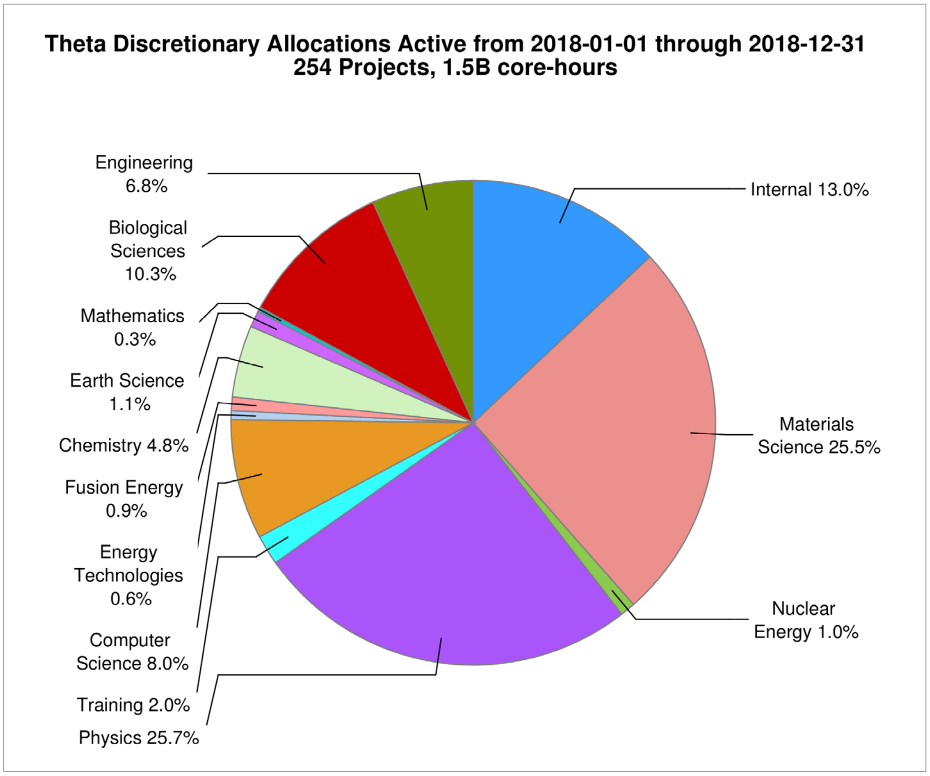

Figure 3.8 Theta CY 2018 DD Allocations by Domain

Conclusion

The ALCF delivered the following core-hours to the allocation programs in CY 2018: 4.6 billion to INCITE, 2.2 billion to ALCC, and 1.4 billion 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?

ALCF Response

Listed below are the innovations and best practices carried out at ALCF during CY 2018. 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

ALCF has undertaken several projects to improve the operations of ALCF and to better respond to user needs.

4.1.1 Continuous Integration

Challenge: Users at ALCF have a need for automation to facilitate development, testing, and release efforts within their project.

Approach: The open source Jenkins solution hosted by ALCF uses configuration techniques and modules to allow ALCF users to integrate with their research projects.

Impact/Status: Continuous integration, CI, is taking shape at ALCF. Available to ALCF users and project groups today is a CI solution using the open-source Jenkins automation server. ALCF users can make use of CI features in Jenkins for their projects such as accessing source code repositories, automating/scripting build and test functionality, use of various triggers to control actions, and submitting jobs to HPC resources.

This approach involved deploying an open-source Jenkins solution that it is currently running a pilot with ALCF users. This solution allows the participating projects to compile and execute code in isolation and utilizes open-source Jenkins plug-ins to provide the connectivity and security requirements of the ALCF. The Jenkins server itself has been deployed to the ALCF's existing VMWare infrastructure, with executors running on the login nodes (Figure 4.1). Jenkins provides connectivity to various repositories, such as Git and SVN, to facilitate the code pulls; the code compilations, builds, and execution occur on the executors. Because one of our goals is to provide separation, each project is creating a service account with the group ownership of the project, and an executor in Jenkins is created per project. The Jenkins executor runs under the service account, and thus the builds occur under the corresponding service account for the project as well. Because we create an executor per project, Jenkins is assigned to launch the executor for the project when the demand for the execution arises. Jenkins provides a history of the builds for the users as well as standard out and standard error.

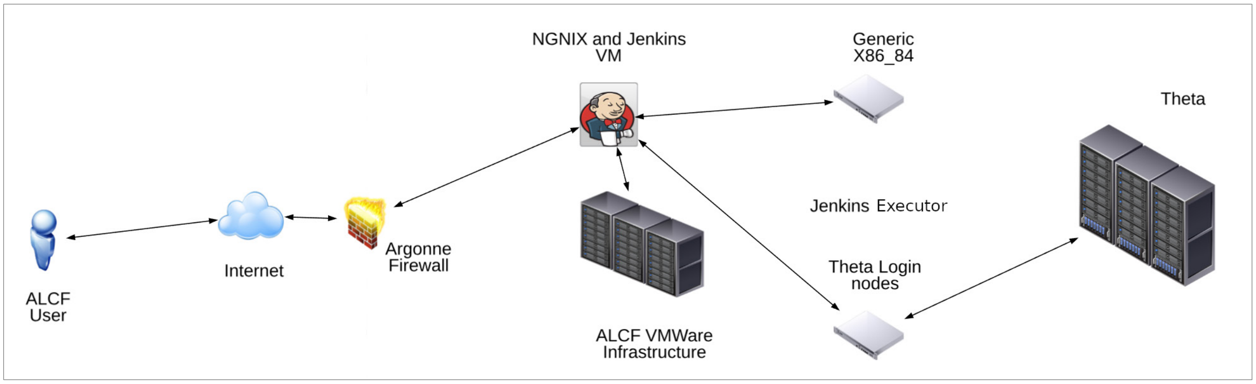

Figure 4.1 High-level Process Flow for the Use of Continuous Integration at ALCF

The larger work flow will include GitLab as part of the ECP project (as discussed in Section 8.3.2). This innovation was independently developed by ALCF staff and was presented at the Cray User Group 2018 (CUG 2018) conference in Stockholm, Sweden.

4.1.2 Web Reports

Challenge: The ALCF reporting infrastructure captures an extensive amount of data from various sources, more than 200 million records per year, and uses graphs and other visual representations to detect anomalies within it. While some data relationships are more difficult to depict in a single graph than others, in order to capture anomalies, it is imperative that all data sets are exposed in one area and are stored in an easily digestible format.

Approach: The ALCF's business intelligence (BI) group worked with various stakeholders, including the operations team, senior management, scheduling committee, and others, to identify the ideal layout for this data to generate the visuals needed for decision making. This effort meant pulling together existing graphs, such as full system outages over time, scheduled/unscheduled availability events over time, and machine usage by program over time, with new graphs, such as the queue depth on the machines covering the same time period. Adding the queue depth information also allowed the team to generate e-mail alerts to key staff groups when the jobs in the queue went below a specified threshold.

Impact/Status: The information presented in these graphs provides insights about many things, including node outages, the impact of an outage on a system's overall availability, the accuracy of time spent draining a machine for new jobs, and whether or not an outage was correctly marked as scheduled ― and can even reveal missing data (Figures 4.2 through 4.5 portray these data for Theta usage for 365, 90, 30, and 7 days, respectively). Individuals are able to quickly extract the information that is most relevant to their work. For example, when the queue depth on the machines falls below a specified threshold, Catalysts can prompt their project teams to take advantage of low queue wait times, thereby keeping resource utilization high.

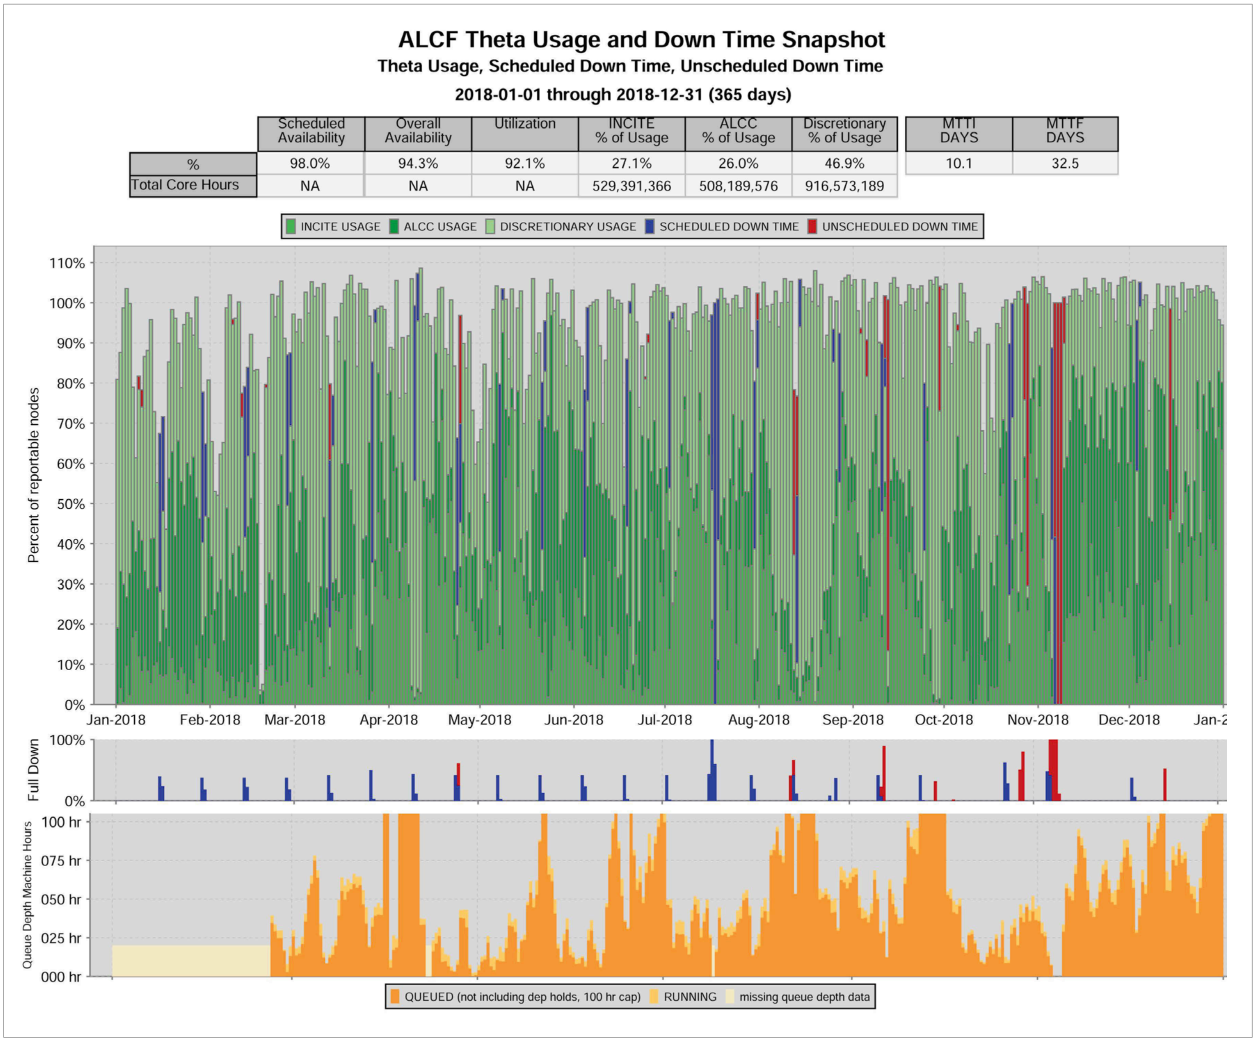

Figure 4.2 Screenshot of Theta Usage, Scheduled Downtime, and Unscheduled Downtime (365 days) *

* Official metric numbers may differ slightly from the values shown on the real-time graphs. The BI Team began collecting queue depth data on February 22, 2018.

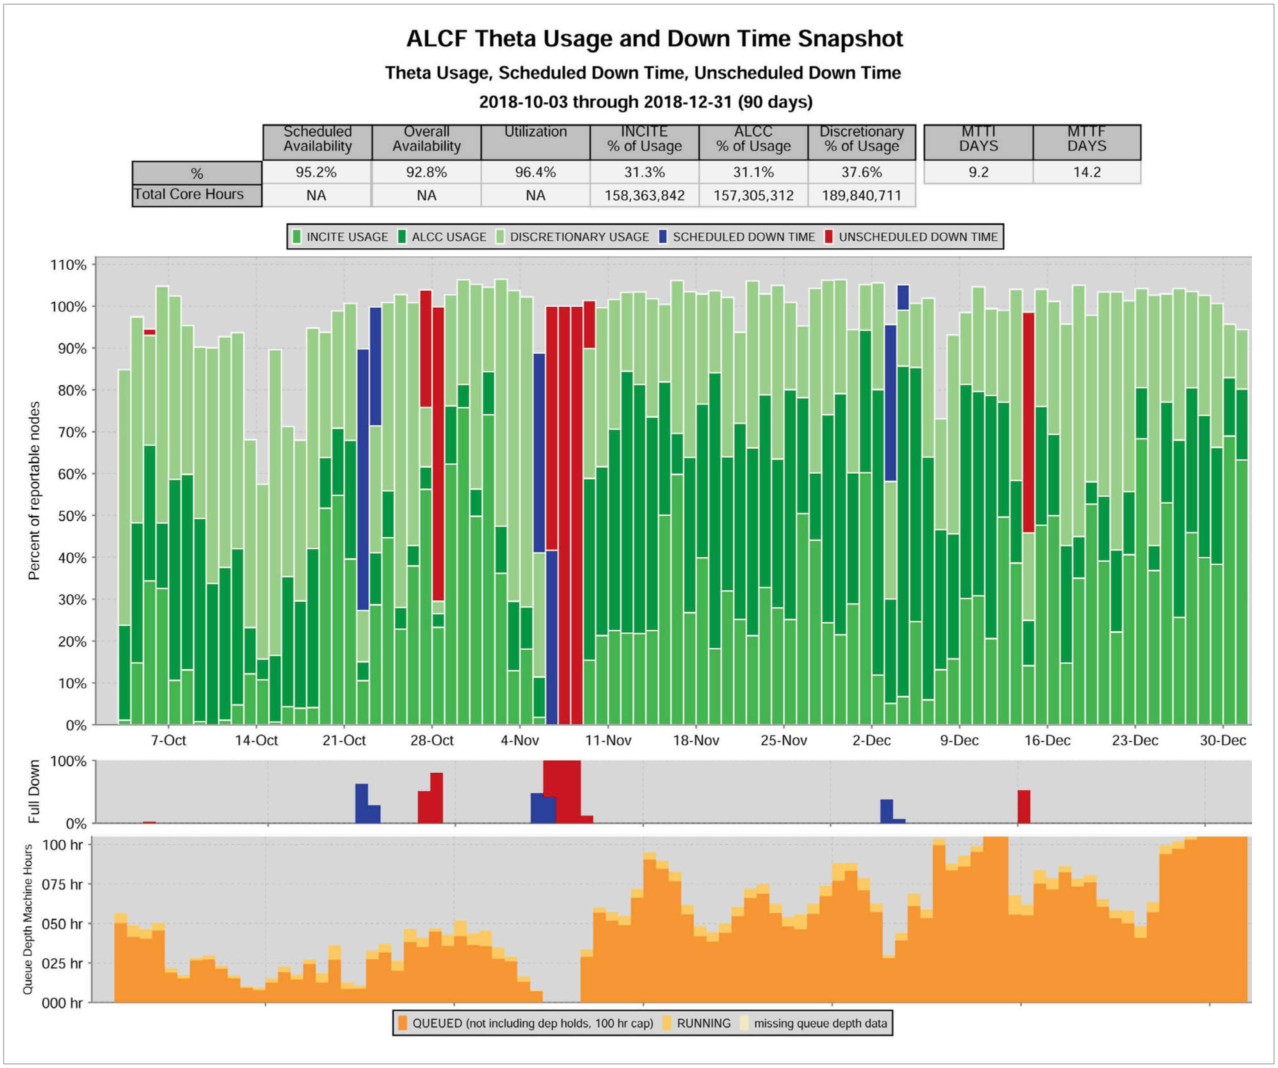

	Scheduled Availability	Overall Availability	Utilization	INCITE % of Usage	ALCC % of Usage	Discretionary % of Usage	MTTI DAYS	MTTF DAYS
%	95.2%	92.8%	96.4%	31.3%	31.1%	37.6%	9.2	14.2
Total Core Hours	NA	NA	NA	158,363,842	157,305,312	189,840,711		

Figure 4.3 Screenshot of Theta usage, Scheduled Downtime, and Unscheduled Downtime (90 days)

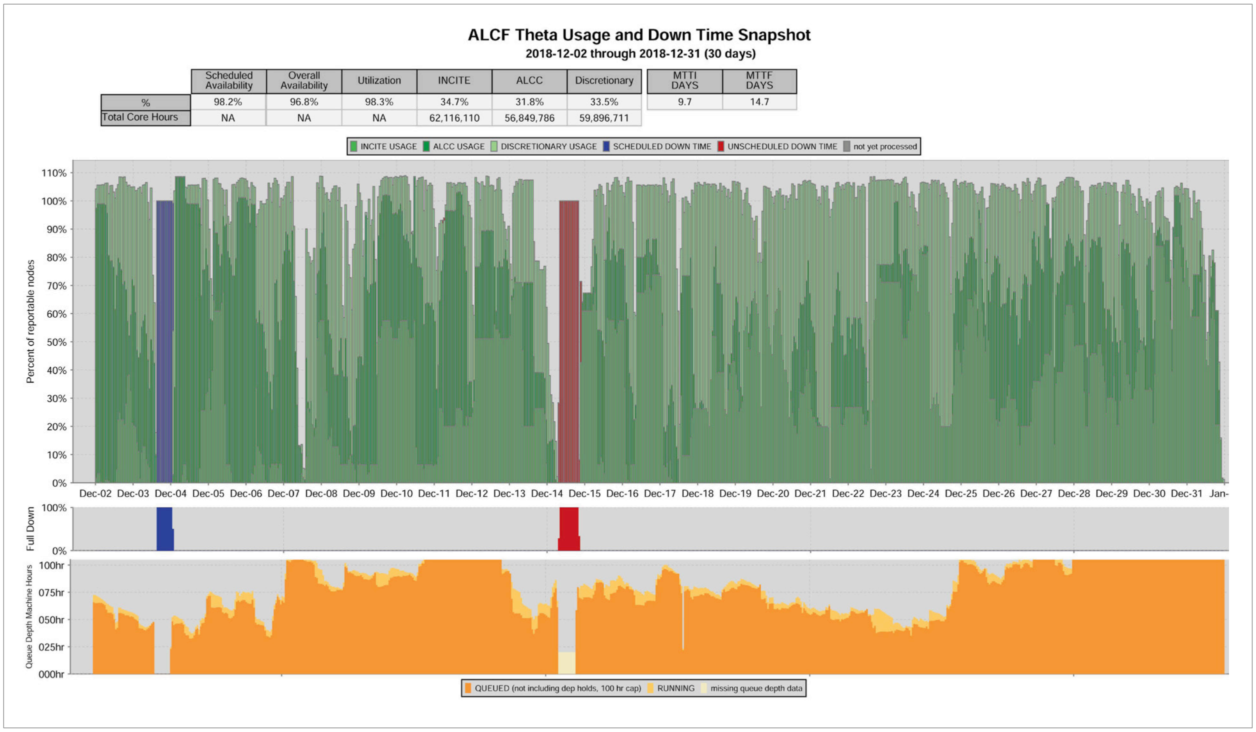

Figure 4.4 Screenshot of Theta Usage, Scheduled Downtime, and Unscheduled Downtime (30 days)

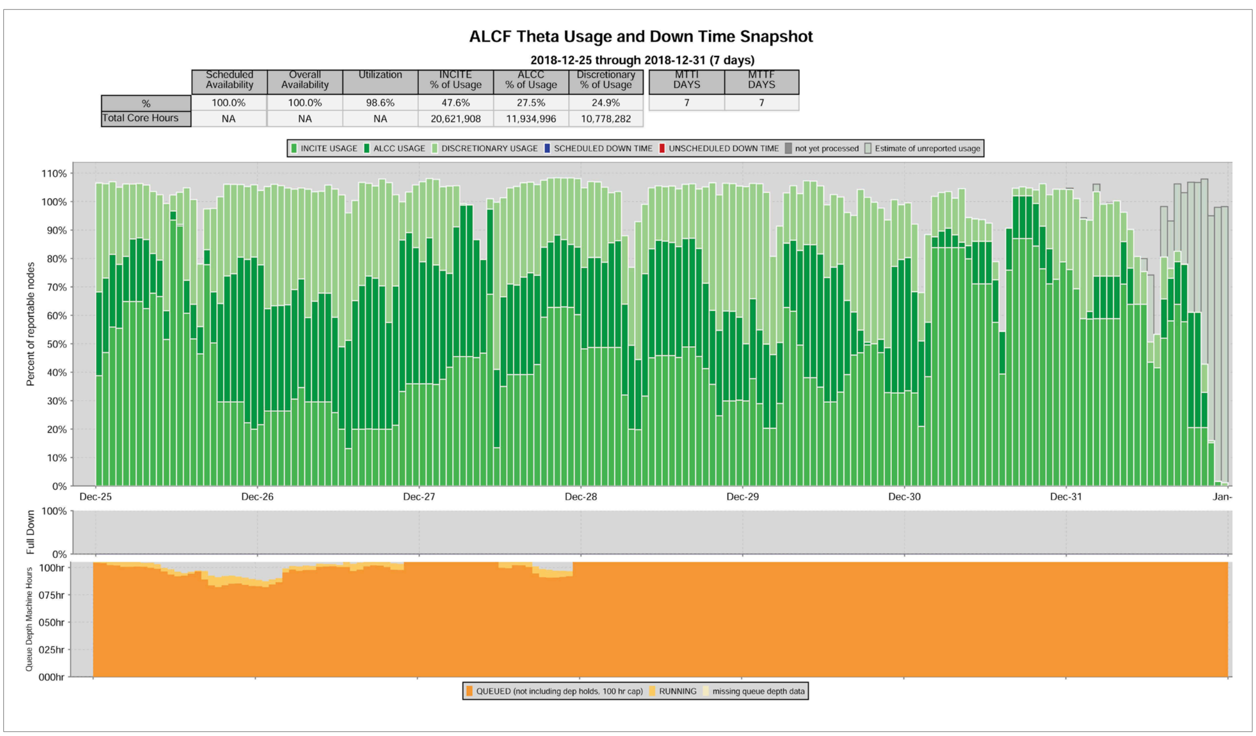

Figure 4.5 Screenshot of Theta Usage, Scheduled Downtime, and Unscheduled Downtime (7 days)

Sample e-mail alert:

ALERT: QUEUE DEPTH FOR MACHINE mira
The Queue Depth machine hours is: 20.46
This alert uses the queue depth plus remaining of running jobs
The Queue Depth does not include Maintenance jobs
The Queue Depth does not include dependency holds

Your alert is 48.00 machine hours

Queue Depth data for Machine: mira
Queued Machine Hours (total): 19.92

Queued Machine Hours (backfill): 9.84

Queued Machine Hours (INCITE): 11.92
Queued Machine Hours (ALCC): 0.73
Queued Machine Hours (Discretionary): 0.00

Remaining Machine Hours of running jobs (total): 6.54

Totals do not include the following:
Queued Maintenance Machine Hours: 6.00
Remaining Maintenance Machine Hours: 0.00

This innovation was independently developed by ALCF staff.

4.1.3 Deployment and Scaling of Containers on Theta

Challenge: Scientific software stacks can have complicated dependencies and build requirements, making it difficult to deploy them on supercomputing systems. The requirements of these software stacks strain the file system performance, especially in the case with large numbers of small, dynamically linked libraries.

Approach: Containers are a single monolithic file, ranging from GBs to >500GB, which contain a full software stack including the system library dependencies. Unlike virtual machines, which emulate the hardware and operating system kernels, containers use the kernel of the local system and its interface to the native hardware. This approach results in better performance of the software.

Results/Status: Singularity is fully supported and working across Theta, Iota and Cooley. For users, Singularity provides the ability to install complex software stacks once in a contained environment; then, by simply copying the container file between computing resources, the application can run. Physicists from the ATLAS experiment at CERN, who study proton collisions, were among the first users at the ALCF to benefit from this newly deployed technology. The complex CERN software and ATLAS analysis framework were installed in a 500-GB container. This advancement was key to achieving linear scaling when running the

ATLAS simulation on Theta. Due to the large size of the software, an installation directly on the standard file system had led to poor performance at large numbers of Theta nodes.

We are currently in the process of testing Singularity 3.0, the next major iteration of the container client. We also have an internal singularity registry, which allows users to store and retrieve their own container images.

ALCF has developed script and best practices to enable the deployment and scaling of containers for science applications. This innovation was independently developed by ALCF staff. This work has been shared with the ECP container working group and presented at ALCF workshops and developer sessions. Documentation of this work appears on the ALCF website.

4.1.4 Theta Monitoring

Challenge: A new monitoring stack for Theta came out of a number of challenges we encountered during the early administration of Theta. The XC40 contains a variety of log sources, and trying to centrally "discover" machine state, operational issues, and trends was difficult, as meaningful data could be spread throughout various log files in different places and in different formats. In addition to this set of circumstances, some data sources were kept in binary files, making analysis a challenge. The alerting capabilities were also limited, and extended only as far as e-mail alerts that fired based on lines in various log files.

Approach: We centralized all of the log data and real-time data from the Cray message bus in an Elastic Stack cluster. This structure allows ALCF staff to visualize all data over time, as well as easily correlate events across the machine from a variety of data sources. All events are parsed to extract location and job identification (ID) information, making it easy to identify and track trends with regard to user behavior, machine state, and other metrics. We can also create more complex alerting logic and send alerts to Slack or e-mail.

Results/Status: Our Elastic Stack solution is being deployed to production, and is already allowing us to create complex and rich visualizations that provide insights into the machine state (Figure 4.6). We also receive Slack alerts whenever an interrupt event or serious fault occurs on the machine (fatal RAS errors, node faults, failed boots, and many other events). We are rapidly adding to the data sources that Elastic Stack ingests. Obtaining real-time performance metrics from Cobalt and ALPS is in progress, as is capturing storage data and gathering metrics from the Theta MOM nodes (a 'MOM' node is a shared service resource that manages job execution).

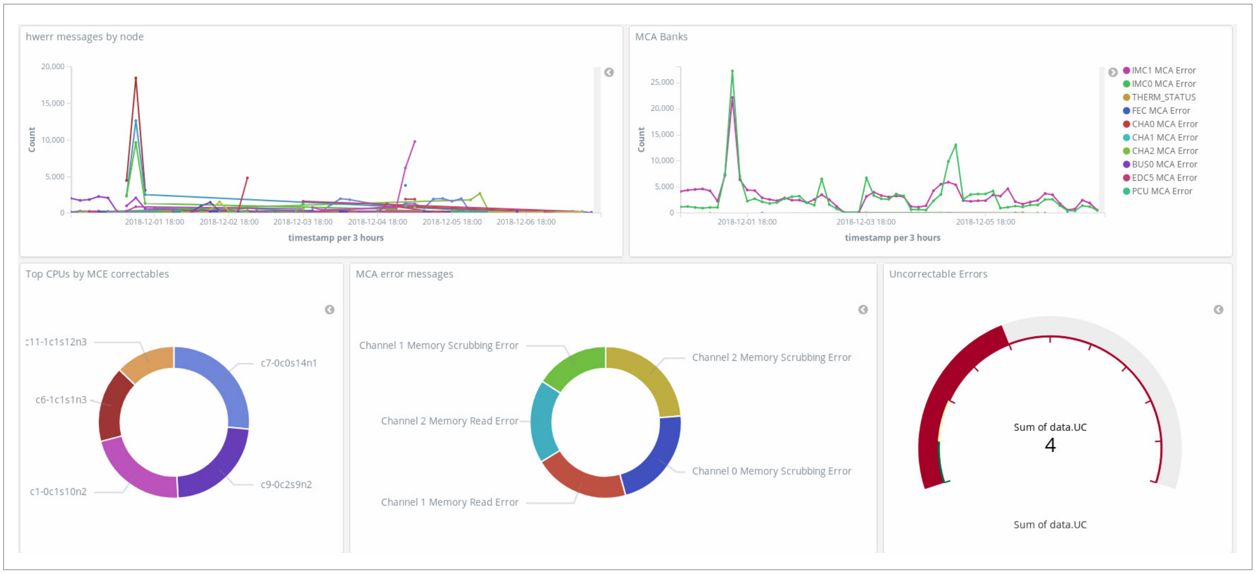

Figure 4.6 Screenshot of a Dashboard Showing CPU Machine Check Errors

This innovation was independently developed by ALCF staff and was presented at the Cray User Group 2018 (CUG 2018) conference in Stockholm, Sweden.

4.1.5 Using Columnar Tables with IBM BLU Acceleration to Improve Data Compression and Storage

Challenge: The ALCF developed a file system monitoring tool for GPFS deployments, called ggiostat, that captures metrics of our 26 PB of GPFS storage. With such a large file system and distributed user base, we needed a way to monitor I/O performance and usage patterns. This goal was accomplished using an in-house tool called Audience. Audience is an application that links GPFS I/O performance with user job data to provide real-time tracking of I/O patterns. It can also query historical data to provide past I/O performance. Originally, the ALCF used a traditional DB2 database to store this time series data. However, with data inserted at the rate of millions of rows per hour, that database ballooned to several billion rows and 3 TB of data. The facility needed a way to maintain the insertion process and query performance while improving data compression to save on space.

Approach: Migrating data from the standard DB2 database to the dynamic, in-memory, columnar table-formatted DB2 BLU addresses the challenge.

Impact/Status: The data have been migrated to DB2 BLU, which has resulted in a data compression level of from 3 TB to about 100 GB. ALCF performance engineers are seeing performance comparable to that of the previous database, and without the use of indexes or structured query language (SQL) tuning. The team has used BLU in other databases to store and analyze data based on these positive results.

Given that the Audience tool was in beta in 2017, here is a report on further progress in the tool usage in 2018:

Most of the work done on the Audience tool in 2018 has involved behind-the-scenes optimization and improvements to the user interface. One such change was adding the option for global filesystem performance statistics and providing the code optimization to make that possible. We also improved the way GPFS data are gathered. Before it was a pull method that required a middleman script to scrape data off the nodes and drop the data into the BLU database. Now the data are directly pushed from the nodes into the BLU database.

This innovation was independently developed by ALCF staff members; however, IBM BLU is used by many organizations, possibly including other DOE user facilities. Colleagues in Argonne's Mathematics and Computer Science (MCS) division are using the ALCF's I/O performance data for research. The ALCF's data compression innovation using DB2 BLU was presented at the 2018 International DB2 Users Group (IDUG) technical conference in Philadelphia, Pennsylvania.

4.1.6 Optimized Builds and Scaling of Distributed TensorFlow and PyTorch Deep Learning Frameworks

Challenge: Deep learning packages are not provided as part of the Cray programming environment, and there is only limited support for these as part of the Cray URIKA stack. However, addressing this gap has a dependency of requiring the shifter container environment, which is not supported on Theta. The Intel Python distribution channel provides for the frameworks; however, these builds are not optimized to leverage the hardware features available on the KNL processor and do not support for scaling to a distributed memory environment. Custom build instructions for the various frameworks from Intel are available; however, the build instructions are not up to date with the newer release versions of the frameworks — a capability requested by our users. In addition, the provided builds faced issues with underlying software environments on Theta, including Cray's MPI and Cray's ALPS affinity settings.

Approach: ALCF staff worked with our vendors, Intel and Cray, to better understand their build process. Next, we developed scripts to build and test the various frameworks optimized to use the software environment as well as the hardware characteristics on Theta. We now provide modules for various TensorFlow, PyTorch, and Horovod versions and a conda environment with support for various deep learning modules as well. We developed a set of benchmark tests to evaluate various software and environment settings, and a set of recommendations for our users, which were presented at the Simulation, Data, and Learning Workshop and the ALCF Computational Performance Workshop.

Results/Status: ALCF users now have the option to choose among multiple builds of various deep learning frameworks to meet the needs of their science application. These builds are now used by several applications, including for ADSP, ECP, ESP, and DD projects.

This innovation was independently developed by ALCF staff. This infrastructure has enabled several publications for science using learning on ALCF systems.

4.2 Research Activities for Future Operations

ALCF innovations are helping prepare for the next-generation systems, both at the ALCF and in the wider HPC community. Additionally, the ALCF tracks research results achieved by the broader research community that are likely to impact the next-generation Facility and its users through its participation in research projects and professional community activities that help drive the state of the art in computer science. This includes work funded by the DOE-sponsored PathForward program, for which the ALCF provides technical representation for the Intel project, regularly reviews the milestones, and provides feedback. Others activities include:

- DOE Storage Systems and Input/Output (SSIO) Workshop
- DOE Computer Graphic Forum (DOECFG)
- IXPUG In Situ Analysis Hackathon
- IEEE Visualization (VIS) (and Large Data Analysis and Visualization (LDAV))
- DOE Performance, Portability and Productivity (PPP) Annual Meeting
- MLPerf HPC forum
- DOE ASCR Scientific Machine Learning Workshop
- NVIDIA Containers working group
- SC'18 Python workshop
- SC'18 including paper on in situ workflows
- Deep Learning for Multi-Messenger Astrophysics workshop

Furthermore, the ALCF is actively engaged in the C++ Standards Committee, the OpenMP Standards Committee, the Open Scalable File Systems (Open SFS) Standards Committee, and the Standard Performance Evaluation Corporation (SPEC) High Performance Group (HPG). The ALCF also engages directly with vendors and researchers through the Low-Level Virtual Machine (LLVM) community, including in-person engagement at the LLVM developers' meetings and vendor-led conference calls (e.g., for Flang, OpenMP, SYCL).

The Facility reports on participation in research projects in Section 1.3.2.1 and professional community activities in Section 8.3.1.

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 CELS computing divisions and run by the ALCF, the JLSE centralizes Argonne's research activities aimed at evaluating future extreme-scale computing systems, technologies, and capabilities. JLSE users 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 by:

- 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 NVIDIA DGX-1 System (Volta graphical processing units [GPUs]); a 3-node Skylake Xeon cluster with 4 V100 SMX2 GPUs in each node; an HPE Comanche with 64 Marvell Arm ThunderX2 processors; an Atos Quantum Learning Machine; a Kubernetes testbed; and a Lustre testbed.

In 2018, the JLSE supported 380 users participating in 60 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: Application 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: Application 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 becoming 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.

PetrelKube: This small Kubernetes testbed is managed with Rancher and hosts backend services supporting web portals running outside of the JLSE and HPC workflows. The testbed utilizes Petrel for its primary data storage, initially staging data over the network and eventually having direct access to Petrel via a POSIX interface.

Neuromorphic: Researchers explore software approaches to power efficient implementation of machine learning algorithms on specialized hardware like Neuromorphic chips.

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.

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

4.3 Best Practices

4.3.1 Separation of Job Failure Analysis and Operational Data

The ALCF's operational data processing system (ODPS) tracks usage, availability, hardware state changes, reservations, jobs, or any event that happens over time that uses a compute resource. This software is used to analyze hardware and job failures on ALCF resources. It is also used to generate availability, usage, and utilization data for Theta. The system was built for processing our metrics quickly and for providing a robust API to the data. Processing a month's worth of data for Theta takes around two minutes.

Job failure analysis (JFA) at the ALCF is performed using an application written internally that presents the failure data on a website. Originally, this application carried out both the processing of data it retrieved from the various data sources and the presenting of data to ALCF staff members to mark and manipulate as appropriate to generate metrics. The JFA site was refactored this year to use the processing power of ODPS instead of processing within the application. This decoupling allowed for quick turnaround times on feature requests on the site.

The JFA Interrupt Tool was refactored into an actor system to simplify debugging and reduce code complexity. This change brought processing pipelines down from hours to seconds. Another effect from implementing the actor system is delivery of extra debugging information on data flow and errors in the system.

4.3.2 Intel NDA Training

ALCF receives highly sensitive information from Intel about the upcoming Aurora system. This information is protected by restricted non-disclosure agreements (NDAs) and must be securely managed by the ALCF. The ALCF developed a process for tracking who is authorized to see the information and a way for people to verify who is allowed access. ALCF also developed required training to ensure that all authorized staff understand their roles and responsibilities.

This page intentionally left blank.

Section 5. Risk Management

Is the Facility effectively managing operational risks?

ALCF Response

The ALCF has clearly demonstrated successful management of both its project risks and operational risks (Section 5.1) in CY 2018. As part of the ALCF's Risk Management Plan (RMP), all risks (proposed, open, and retired) are tracked, along with their triggers and mitigations (proposed, in progress, and completed), in a risk register managed by risk managers. All risk ratings in this report are post-mitigation ratings. The 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 implementation of austerity measures as necessary; and (2) Staff Recruitment Challenges, which is managed by ongoing recruiting and re-tasking 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, as well as the major risks that will be tracked in CY 2019 (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:

- *A brief overview of the risk management process employed by the Facility, including the cycle for identifying, mitigating, and retiring risks;*
- *A brief summary of the key risks and their mitigations, including:*
  - *The 3–5 most important operational risks for the review year;*
  - *Any significant new operational risks since the previous review;*
  - *The operational risks that were retired during the review year;*
  - *The major risks that will be tracked in the next year; and*
  - *For the risk events that occurred, how the Risk Management Plan was implemented and an assessment of the mitigations' effectiveness.*

5.1 Risk Management Process Overview

The ALCF uses the documented risk management processes, first implemented in June 2006 and outlined in its RMP, for both operational and project risk management. This RMP is a strategic plan that is annually reviewed, and updated as needed throughout the year, to reflect changes, to incorporate new risk management techniques as they are adopted, and to incorporate best practices from other facilities. Risk management is part of ALCF's culture, and RMP processes are part of normal operations and all projects, such as the ALCF-3 project launched in CY 2013.

Risk management is an iterative process that includes identifying and analyzing risks, performing response planning, and monitoring and controlling risks as shown in Figure 5.1.

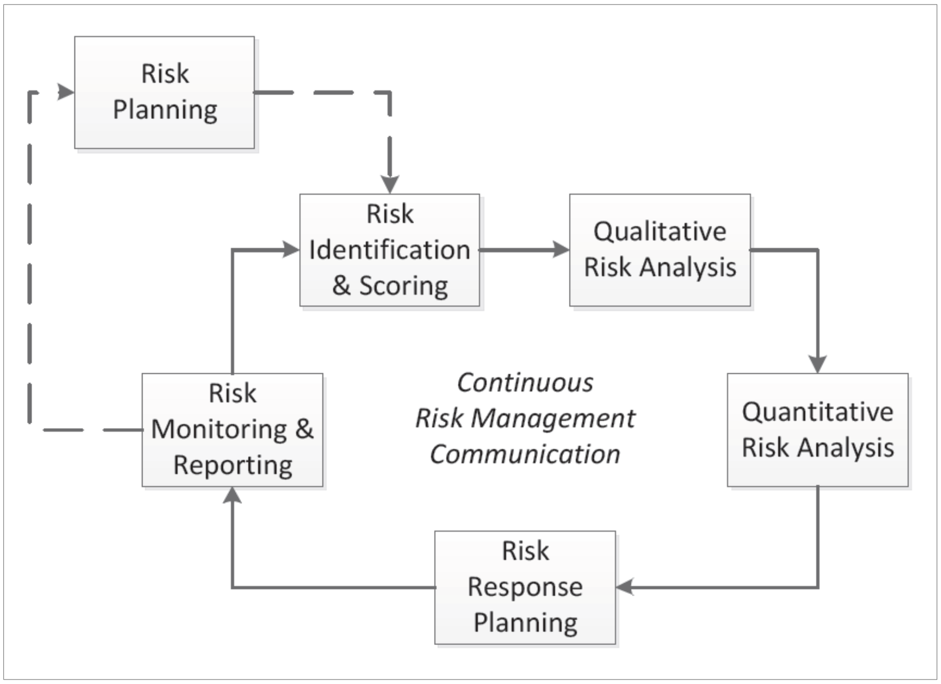

Figure 5.1 Illustration of the Iterative Risk Management Process from the ALCF's Risk Management Plan

The ALCF risk management process consists of the following steps, which are performed on a continuous basis in all normal operations and in all ALCF projects:

1. Plan, implement, and revise the RMP.
2. Identify threats and opportunities to the cost, schedule, and technical objectives.
3. Analyze the impact of identified threats and opportunities to the cost, schedule, and technical baselines; and develop risk management strategies to manage the risks.
4. Monitor risks, mitigation plans, and management reserve and contingency until the risks are retired or a project is closed.

The objective of this process is to identify potential threats and opportunities as early as possible so that the most critical risks can be assessed, the triggers effectively monitored, and the amount of management reserve/contingency needed to moderate the risks determined.

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 Risk Review Board

The ALCF employs a five-person Risk Review Board with representatives from senior management, the operations team, the science team, industry outreach, and the financial services team to serve in an advisory capacity to ALCF management. The board meets as needed and offers recommendations regarding steady-state risk management issues. The RMP is consulted at all risk meetings. At each meeting, the board:

- Reviews proposed new risks and makes recommendations on adding a proposed risk to the steady-state 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 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 as a result of 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.
- Review and identify any risks to retire.
- Reviews the risks encountered in the past 18 months to discuss potential actions.
- Discusses risks encountered at other facilities and identifies any that apply to the ALCF.

5.1.2 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 and individual risk meetings and the Risk Review Board meetings, the ALCF holds many informal risk discussions. Risks are identified and evaluated, and mitigation actions are 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 developed in consultation with safety and subject matter experts. JHQs are used for all staff and all contractors and cover all work. During planning meetings for any 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 machines, 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.

5.1.3 Continuation of the ALCF-3 Project

The project to procure and deploy the ALCF's next supercomputer, known as ALCF-3, continued in CY 2018. 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.

5.2 Major Risks Tracked for the Review Year

The ALCF has been operating both Mira and Theta during CY 2018, and planned the growth of both the staff and the budget in order to bring the facility to full strength. As such, the ALCF continues to monitor a large number of major risks for the facility. No major risks were retired during CY 2018.

Four major operations risks were tracked for CY 2018, **two** with a risk rating of High, **one** with a risk rating as Moderate, and **one** with a risk rating of Low. Of these, three were encountered and managed. The **four** 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 2018

ID	Title	Encountered	Rating	Notes
1059	Funding/Budget Shortfalls	Yes	High	The ALCF worked with the program office to plan a budget for handling the impact of a Continuing Resolution in FY 2018, 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.
25	Staff Recruitment Challenges	Yes	High	The ALCF added 14 new hires overall this year, plus 2 internal reclassifications. The ALCF continues to have staff available who can be re-tasked as needed. With difficulty competing with industry for new hires, staff hiring remains a concern.
1049	Staff Retention	Yes	Moderate	The ALCF lost **3** staff members during CY 2018. 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.
1091	Injury to Workers/Overall Safety of the Division	No	Low	

5.3 Risks Encountered in the Review Year and Their Mitigations

The **six** risks encountered during CY 2018 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 Low, and 3 Very Low.

5.3.1 Funding/Budget Shortfalls

1059: Funding/Budget Shortfalls	
Risk Owner	Michael E. 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 the 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. The 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 2018; 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, the ALCF prepared for a full-year Continuing Resolution and reduced budget scenarios. To ensure that adequate funds would be available to operate Mira and Theta and to continue ALCF-3 activities, ALCF maintained moderate austerity measures to provide maximum flexibility for the coming fiscal year.

The FY 2019 budget was approved prior to the start of the fiscal year. The ALCF will continue to work closely with DOE-ASCR to monitor the budget information and to prioritize spending.

5.3.2 Staffing Recruitment and Retention Challenges

25: Staff Recruitment Challenges	
Risk Owner	Michael E. 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. Re-task 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 E. 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 the 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 2018 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 full-time regular and two long-term temporary. **Seven** new full-time regular staff (five external hires, two internal hires from other Argonne divisions) and **eight** long-term temporary staff (six external hires, two reclassifications from postdoc/predoc) were added during CY 2018, for a net **gain of 12** ALCF staff members for the year. While the 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.

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

By carefully and judiciously managing both risks, the 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 the ALCF.

5.3.3 System Stability Issues Due to Upgrades

1056: System Stability Issues Due to Upgrades	
Risk Owner	John P. Reddy
Probability	Very Low
Impact	Cost: Moderate; Technical Scope: Very Low
Risk Rating	Very Low
Primary Management Strategies	Perform upgrades on non-critical systems first when feasible. Have a rollback plan in place. Monitor performance closely following upgrade. Work with the vendor to understand the upgrade(s) and the quality control processes. Deep test on Test and Development Systems.
Triggers	Planned system upgrades. System instability observed following system upgrades.

Description

Applying a security patch on Theta triggered a bug in other software on the machine.

Evaluation

In early February, we applied Cray-provided security patches to address the highly publicized Spectre and Meltdown vulnerabilities. This action triggered a bug where the VTune software would cause nodes to crash on Theta.

Management

VTune was disabled. Intel provided patches for VTune, which were deployed during the next scheduled preventative maintenance.

5.3.4 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 help ensure 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.

Event 1:

Description

On February 5, 2018, water temperatures in the cooling loop increased to higher than normal. This temperature increase did not cause any problems for Mira or Theta.

Evaluation

Cooling loop water temperatures varied more than expected during the month of February. We noticed the increased temperatures on the Theta and Mira sensors. Although the temperatures varied, they did not reach a critical point on either system.

Management

The strainers were cleaned on February 26, 2018, which improved the situation. ALCF staff continued to closely monitor the temperature and continued working with infrastructure services and building management to preempt any further issues.

Event 2:

Description

On a planned maintenance day, an ALCF staff member was working with building and facilities staff to troubleshoot heat exchanger issues in the Chilled Water Plant.

Evaluation

Infrastructure services and facilities staff switched over to a new heat exchanger but neglected to first clear the system of air, which caused temperatures to spike until the air was bled out. Theta nodes shut down temporarily in response to the high temperatures.

Management

Theta returned to normal operation once the heat exchanger was functioning normally.

Event 3:

Description
A 6-foot-long water line separated in Argonne's 200 Area water cooling loop supplied by the Chilled Water Plant resulting in a catastrophic loss of return water. The plant was unable to provide chilled water for nearly 4 hours, causing the water temperatures to rise in the Mira loop along with the air handling units (temporarily supplied by Area 200 water).

Evaluation
High air and water temperatures caused various resources in the data center to shut themselves off. The water line was repaired and chilled water operations resumed. The ALCF received notification of the incident upon completion of the fix.

Management
Argonne Infrastructure Services has instituted corrective actions to ensure that known information is relayed in a timely manner going forward.

5.3.5 Users Gain Access to Data That They Should Not Have Permission to Review

1046: Users Gain Access to Data That They Should Not Have Permission to Review	
Risk Owner	John P. Reddy
Probability	Very Low
Impact	Cost: Very Low; Technical Scope: Very Low
Risk Rating	Very Low
Primary Management Strategies	
Triggers	Scan of files shows bad permissions on files. User notifies support of an issue.

Event 1:

Description
While conducting a proactive internal audit, ALCF staff discovered that some of the users on the systems had their home directories set to world readable.

Evaluation
Some of these home directories contained data that users are not authorized to access.

Management
To mitigate encountering this risk in the future we created an internal policy detailing how home directories should be set up and secured. We developed an internal tool to automate scanning permissions and to send notification of home directories that have not been properly locked down to the cyber security representatives.

Event 2:

Description

On June 17, 2018, a Sunday, a file transfer protocol (FTP) server in Argonne's Computing, Environment and Life Sciences (CELS) directorate restarted and began serving information it was not expected or supposed to be sharing.

Evaluation

Data exposed included two items of note for the ALCF: the source for an old instance of the ALCF's user management website, including a password to the user management database; and information for a test instance of the ALCF's user management website, including a password to a copy of the user management database. Both databases contain personally identifiable information (PII).

Management

Access to this FTP site was shut down the following day upon discovery. Affected passwords were changed within hours of notification of the exposure. The old copy of the web page was removed. The test instance of the database has been deleted. The affected databases were not, and are not, accessible from off-site. ALCF Operations is auditing all data stores, databases, and software repositories to identify other potential avenues of exposure.

Event 3:

Description

On September 13, 2018, during a routine auditing process, ALCF's cyber security program representative discovered a misconfiguration in services provided by the CELS systems group.

Evaluation

The TRAC service was exposing some subversion repositories that are supposed to be kept under strict access control.

Management

Access via TRAC was immediately terminated, configurations were reviewed and corrected, and then TRAC was brought back online. All customers of the service were contacted and informed of the possible exposure.

5.3.6 Facility Power Interruptions

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

Event 1:

Description
A breaker tripped on a substation that feeds one of the Theta wall panels. This dropped power to six cabinets and a blower cabinet, which cascaded into all cabinets in row 1 going down. This event happened on September 11, 2018, and again on September 29, 2018.

Evaluation
The trip unit in the breaker was assessed as likely defective. After the first incident, the facility manager began the process of ordering a replacement breaker and a spare. As the breaker passed tests after the first incident, it was reset. After the second occurrence, the replacement order was expedited, and the facility manager readied an interim fix if the problem were to reoccur before the replacement arrived.

Management
Substation was brought back online after tests and subsequently Theta was brought back online after power was restored to the data center. Future work included replacing the faulty breaker.

Event 2:

Description
We conducted a planned outage so that facility management could work on three power-related issues. The only one affecting Theta was replacing the faulty breaker. For more details about this incident, please refer to Section 2.1.2 — September outages.

Evaluation
Routine maintenance work was performed to address earlier power-related issues.

Management
There are no proposed actions. One of the items addressed was installing a new temporary breaker to provide a stopgap to the September breaker problems that took out Theta twice.

Event 3:

Description

The substation feeding Theta was turned off after testing on preventative maintenance day indicated possible low resistivity.

Evaluation

Offsite testing was performed, and after a few days, we were informed that it was safe to re-energize the substation.

Management

Once we were notified of the safety of the substation, we brought Theta back online. For more details about this incident, please refer to Section 2.1.2 – November outage.

5.4 Retired Risks

There were no risks retired in CY 2018.

5.5 New and Recharacterized Risks since the Last Review

There are no new and recharacterized risks since the last review.

5.6 Projected Major Operating Risks for the Next Year

Table 5.2 lists the current top operating risks projected for CY 2019, 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 5.2 Projected Operating Risks for CY 2019

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. Re-task 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

The 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. 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.

Section 6. Environment, Safety, and Health

(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 Environment, Safety, and Health measures?

ALCF Response

The ALCF has implemented appropriate Environmental, Safety, and Health (ES&H) measures and exhibits a culture of continual improvement in ES&H practices to benefit staff, users, the public, and the environment. The ALCF has maintained an exemplary safety record since its inception in 2006 and has never experienced a lost time incident.

A major focus of the Argonne Site Office Performance Evaluation and Measurement Plan for FY 2018 that impacts the Facility is the implementation and continual improvement of Argonne's electrical safety program. The laboratory initially revamped its electrical safety program in late 2017 and has continued to put into effect additional procedural improvements since that time. Elements of the revamped program implemented to date have focused on more detailed work planning, the effective utilization of Lab subject matter experts early in the work planning process, and a higher rigor qualification process for electrical workers. As the ALCF had implemented its own electrical program, which was designed specifically for our scope of work and workers, prior to this time, activities within the Facility have centered around collaborative efforts with the Argonne Authority Having Jurisdiction (AHJ) and directorate safety staff to transition to the Lab program and to vigorously vet both our work procedures and electrical workers. Conceptually, this was not a significant change, and primarily consisted of switching from ALCF-designed and delivered training program classes and requirements to Argonne's lab-wide training and requirements, as well as updating our fundamental work planning guidance documentation to reflect the new Lab program and associated structure and goals.

The ALCF has overhauled our local work planning and control manual to reflect the electrical program changes and to incorporate approaches from the DOE-HDBK-1211-2014: Activity-Level Work Planning and Control Implementation. The division now employs the Laboratory's web-based application Aware for hazard identification, risk mitigation controls, and worker authorization. A formal "skill of the worker" document was developed collaboratively with the Data Center Facility Manager, a union electrician, to cover routine low-rigor, non-electrical worker system administrator tasks. Formal task-based work control documents 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 due to water treatment chemicals, weight, and potential damage to hardware), as well as medium-voltage electrical maintenance. Management actively monitors work scope, updates hazard analyses as they are identified, and revises work planning and control documents accordingly, including for emergency work or when there is an unexpected change to previously planned work. Worker health and safety documentation that cannot be housed within Aware is retained either within the Laboratory's Training Management System or the Laboratory's official cloud resource Box.

A formal Management Assessment was scheduled and conducted right after the start of the calendar year to ensure a smooth and effective transition to the Laboratory's qualified electrical worker program. Minor improvement opportunities were identified and a single issue, all of which were immediately corrected. The ALCF Operations team continues to implement a strong Integrated Safety Management cycle by opening all operations meetings with a safety discussion and feedback that is immediately translated into improved operations as identified. All work control documents incorporate ES&H Coordinator review, as well as input from subject matter experts, including the Data Center Facility Manager and Laboratory AHJ. Pause work is encouraged by management and liberally used by staff to ensure that no unidentified hazards remain, either from planned or co-located work within the Data Center. All major planned operations are shared with the Data Center Facility Manager in advance to strengthen the Facility's safety program; and all non-normal events, no matter the magnitude, are investigated.

Section 7. Cyber Security

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

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

ALCF Response

The ALCF exhibits a culture of continual improvement in cyber security practices and has improved its overall cyber security posture. In CY 2018, the ALCF worked with Argonne's Cyber Operations, Analysis, and Research (COAR) team to implement multi-factor authentication (MFA) for Argonne. Authentication using DOE-issued Homeland Security Presidential Directive (HSPD)-12 badges is required for all privileged users to access ALCF resources. All privileged ALCF staff hold Level 4 HSPD-12 cards. A review of the ALCF cyber security plan was also completed.

In 2018, there were **0** cyber security incidents on ALCF-managed systems. The ALCF's cyber security personnel take a proactive approach to problem management. Here are examples of the proactive steps taken by our security personnel in the past year:

1. Conducting reviews to ensure that users have correct access privileges.
2. Conducting reviews of data storage protocols and security measures.
3. Educating users and staff on how to secure passwords in applications to prevent accidental exposure.
4. Educating developers on secure coding best practices via internal discussions/reviews and external training courses.
5. Updating the National Institute of Standards and Technology (NIST) Certification package, including NIST 800-18, 800-30, 800-34, and 800-53 compliance documents.
6. Evaluating Duo as a new MFA method in a JLSE effort.
7. Archiving and deleting old data.
8. Developing a password scanning application to identify passwords stored in Subversion (SVN)/Git repositories so they can be scrubbed.
9. Verifying that passwords are being rotated on a regular basis.
10. Monitoring new vulnerabilities as they are released and evaluating whether they affect our systems.
11. Conducting penetration testing of internal- and external-facing web applications to evaluate what improvements can be made and how we can better protect our users.
12. Developing tools to verify permissions on home directories following ALCF best practices.

The ALCF has a valid cyber security plan in place. All security vulnerabilities are thoroughly investigated and fixed. In CY 2018, none of the security issues that were found were known to be exploited. Here are several examples of security issues that were detected in 2018, and their ensuing mitigations:

1. Detection of a vulnerability in a user management interface that would allow an SQL injection to access user details. We immediately reworked our processes to disable the affected features. Note: This interface did not contain any PII.
2. Detection of a vulnerability in a user management system that allowed brute forcing of an e-mail verification code that allowed access to user details. This was fixed by working with the developers to patch the code and redeploy. Note: This system did not contain any PII.
3. Detection of SVN repositories hosted by another Argonne division exposing ALCF SVN repositories because of erroneous permissions settings. We immediately shut off the service and kept it offline until we could verify that all repositories were properly secured and only the authorized users had access.
4. On an internal ALCF network, a web interface for our configuration management system containing system configuration details was exposed to internal users. The web interface was updated to require MFA and access restrictions were added. Note: There was no indication that any sensitive configuration data were improperly accessed.
5. Detection of a vulnerability in the ALCF's internal GitLab server for system administrators that could have enabled unauthorized access to users and the ability to take over another user's account. We removed this vulnerability from the system.
6. A security audit of an internal allocation approval application turned up a number of vulnerabilities that could be used to escalate privileges for authenticated (known) users. We worked with the developer to address all security issues found in the audit and provided guidance on how to further secure the application.
7. A separate Argonne division that hosts services for the ALCF had an FTP server that began serving the root of the file system. The provider detected that data were being downloaded by an external group. This exposed code to one of ALCF's applications, and the code contained a password that had access to PII. Access to the database was restricted to five ALCF-managed hosts that are accessible only to system administrators. An investigation confirmed that no one had used the password to access the data. The password was rotated immediately to further eliminate the threat.

The ALCF will continue to take a proactive stance on security issues that will further strengthen our cyber security posture. Some plans for increasing ALCF security in the future are as follows:

1. Implementing HSTS (HTTP Strict Transport Security) and HTTPS (Hypertext Transfer Protocol Secure) on all internal and external services.
2. Minimizing the attack surface by retiring services and data that are no longer needed.
3. Verifying that strong encryption is used everywhere in the environment and that plain text protocols are not used for production needs.
4. Improving real-time log analysis techniques.

Argonne's Cyber Security Program Office (CSPO) started a Division Site Assist Visit (DSAV) program to help ensure that all Argonne divisions are following cyber security best practices and can obtain any help that is needed. For the 2018 DSAV, the CSPO awarded ALCF perfect scores in 7 of the 9 categories: authentication and authorization; configuration management; firewall/intrusion detection system; logging and analysis; remote access; training and education; and vulnerability scanning. The ALCF has taken the guidance of the CSPO in the other two areas (incident response, and policies and procedures) and has made improvements to help increase the cyber position of the division.

In October 2018, Argonne was visited by the DOE EA-21 Cyber Security Audit team. While we are still awaiting the final evaluation results, the preliminary results did not turn up any findings in the ALCF. The CSPO indicates that the final audit results will not turn up any new findings.

The ALCF has a valid Authority to Operate (ATO). The Argonne 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. Strategic Results

(a) Are the methods and processes for monitoring scientific accomplishments effective?

(b) Has the Facility demonstrated effective engagements with strategic stakeholders (i.e., beyond the user population)?

(c) Is the Facility operating in a manner that enables delivery of facility mission and Department of Energy mission including maintaining a vibrant US effort in science and engineering?

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:

- Scientific Accomplishments;
- Scientific Highlights;
- Stakeholder Engagement and Outreach.

8.1 Scientific Accomplishments

The process for determining a scientific highlight varies, but typically originates from a member of the Catalyst team and is based on an outcome documented in a quarterly report. Additional sources of potential highlights include technical communications with the PI or co-PI of a project; significant findings reported in a high-impact publication or conference presentation; and a Catalyst's own involvement in a publication. Other measures used to gather user publications outside of the quarterly reports include sending e-mail reminders to users to verify their publications listed on the ALCF website, and performing Google Scholar and Crossref searches. Determination and coordination of highlights are made during the monthly meeting of the ALCF's applications team, which is composed of Catalysts and members of the performance engineering team, and done in consultation with ALCF's director of science.

The ALCF has been collecting ORCIDs (Open Research and Contributor IDs) from its users for the past couple of years. We are investigating ways of using this method to collect more user publication data. Providing us with an ORCID has been completely optional, and we have collected this data for approximately one-third of the users active in a given year.

Table 8.1 shows the breakdown of refereed publications based, at least in part, on the use of ALCF resources, and highlights those appearing in major journals and proceedings. These include three publications in *Nature*, one in *Nature Chemistry*, and eight in *Nature Communications* (combined in the Nature journals category in the table below); the journal *Science*, *Proceedings of the National Academy of Sciences (PNAS)*, *Physical Review Letters,* and the proceedings of the *2018 International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. Table 8.2 shows updated publication counts from prior years and are based on new information after the prior year's OAR deadline.

Table 8.1 Summary of Users' Peer-Reviewed Publications in CY 2018

Nature Journals	Science	PNAS	Physical Review Letters	SC	Total 2018 Publications
12	3	3	13	2	276

Table 8.2 Summary of Users' Peer-Reviewed Publications for 5-year Moving Window

OAR Year	CY 2014	CY 2015	CY 2016	CY 2017	CY 2018
Total Publications	193	181	199	225	276

8.2 Scientific Highlights

Scientific highlights are short narratives that illustrate the user facility's contribution to advancing DOE strategic goals. Each project highlight includes a figure and a bar graph showing time allocated and time used. The first number in the graph title is the allocation total and the second is what the project used. The individual bars represent the percentage of time used on the fraction of the machine shown below the bar, which are "no capability", "low capability", "high capability" from left to right.

8.2.1 Multiscale Physics of Ablative Rayleigh-Taylor Instability

The Science

The Rayleigh-Taylor instability (RTI) occurs when a heavy fluid is accelerated against a light fluid, arising in many natural and engineering systems such as the oceanic meridional overturning circulation, jet-driven lobes in intergalactic clusters, supernova explosions, and inertial confinement fusion. A major difficulty in the numerical modeling of such flow systems exhibiting RTI is the vast range of scales involved, all of which are dynamically coupled given the highly nonlinear nature of the flow.

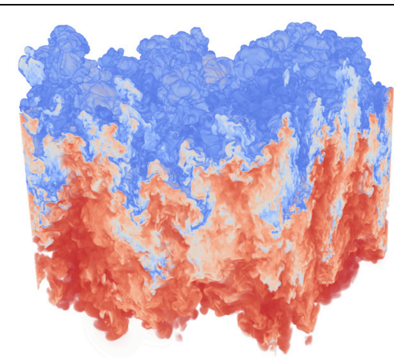

Figure 8.1 Density snapshot of a grand-challenge, fully compressible Rayleigh-Taylor simulation at a resolution of 1,024^2 × 2,048.

(Image courtesy of Dongxiao Zhao, University of Rochester.)

The Impact

The research team has found that any length scale in ablative RTI can be destabilized — contrary to current modeling practices, which assume that small length scales are stable and can be neglected. Their findings have important ramifications for modeling efforts in implosion physics, astrophysics, and combustion science.

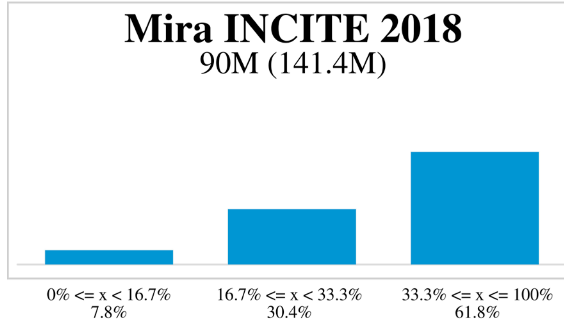

Summary

Ablation is mass evaporation due to a heat source such as the laser in laser-driven plasmas, and it occurs in many other RTI flows, such as in supernova explosions and in molecular clouds. The researchers have found that any length scale in ablative RTI can be destabilized if the initial perturbation is sufficiently large to be in the nonlinear regime from the onset, that is, without passing through the linear regime. This phenomenon is especially important for the small length scales: although often neglected, they were found to be more efficient at driving mixing. Figure 8.1 shows a density field in a compressible Rayleigh-Taylor instability where a heavy fluid (blue) is placed on top of a light fluid (red). The instability is triggered at the interface between the two fluids by means of small-scale random perturbations. Using the vast dynamic range afforded by their high-resolution simulations, the researchers were also able to provide numerical evidence in support of their mathematical results showing the optimal way to analyze length scales in highly nonlinear flows with significant density variations. Results based on this optimal scale decomposition have practical modeling implications for showing conditions under which microphysical effects can be neglected in simulations. ALCF contributions include staff support, assistance with reservation requests, and special queue accommodation.

Contact

Hussein Aluie
University of Rochester
hussein@rochester.edu

Publications

Zhang, H., R. Betti, V. Gopalaswamy, R. Yan, and H. Aluie. “Nonlinear Excitation of the Ablative Rayleigh-Taylor Instability for All Wave Numbers.” *Phys. Rev. E* **97**(1), 011203 (2018). [https://doi.org/10.1103/PhysRevE.97.011203]

Zhao, D., and H. Aluie. “The Inviscid Criterion for Decomposing Scales.” *Phys. Rev. Fluids* **3**, 054603 (2018).

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR

August 2018

8.2.2 Toward Breakthroughs in Protein Structure and Design

The Science

Synthetic mini-protein molecules (peptides) have potentially revolutionary applications in medicine. Thus, the computer-guided design of synthetic peptides (built from synthetic and natural amino acids) with targeted properties is a grand challenge in structural biology.

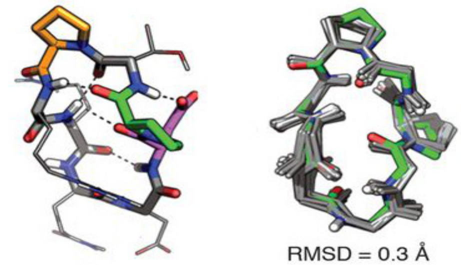

Figure 8.2 Left: a synthetic peptide made of a mixture of natural L-amino acids and mirror-image D-amino acids. Key amino acid residues required for rigid structure are shown in orange, green, and pink. Right: overlay of an experimentally determined nuclear magnetic resonance (NMR) ensemble (grey) on original computational design (green). This peptide has been experimentally confirmed to have a rigid conformation matching the designed conformation. The structures differ by a root mean square difference (RMSD) value of 0.3 Angstroms.

(Image: Hosseinzadeh, P., *et al*., *Science* **358**, 1461–1466 [2017].)

The Impact

For many years, the Baker group has been developing new energy functions and sampling algorithms for peptide structure prediction. This most recent work enabled the design of cyclic peptides that can fold into well-defined shapes with target functions.

Summary

As is the case for protein design, the key to designing a peptide drug is creating a molecule with a rigid structure (Figure 8.2) that perfectly complements the shape and charge of a binding pocket in a target protein (e.g., a pathogen enzyme to be inhibited). However, designing peptides is more challenging than designing proteins, for several reasons. Mira has allowed the researchers to study mixed-chirality peptide macrocycles such as cyclosporine, which are among the most potent therapeutics identified to date. These results validate the algorithms and energy functions developed in recent years for computer-guided design of peptides. Of the 200 folded, stable structures predicted, 12 structures were examined via nuclear magnetic resonance (NMR); of these 12, 9 structures were correctly predicted by computational simulation.

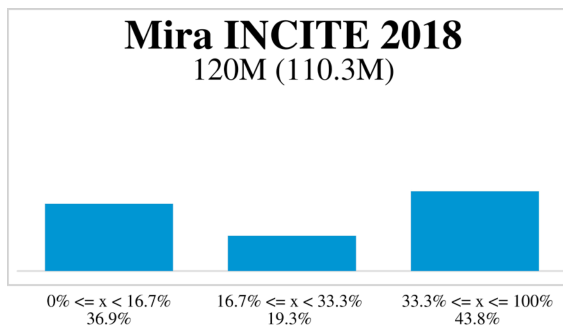

Contact

David Baker
University of Washington
dabaker@uw.edu

Publication

Hosseinzadeh, P., *et al*. “Comprehensive Computational Design of Ordered Peptide Macrocycles.” *Science* **358**, 1461–1466 (2017).

Highlight Categories
Performer/Facility: ASCR-ALCF

Date Submitted to ASCR
March 2018

8.2.3 The Nature of Interlayer Binding and Stacking of Carbon Allotropes

The Science

Graphyne is a theoretical allotrope of carbon that is similar to graphene but has yet to be synthesized, although its existence has been conjectured since the early 1960s. The band structure of graphyne is predicted to contain Dirac cones on its double- and triple-bonded carbon atoms. Because materials with Dirac cones can lead to the manufacture of significantly faster transistors and electronic components, it has been the subject of multiple *ab initio* studies ever since.

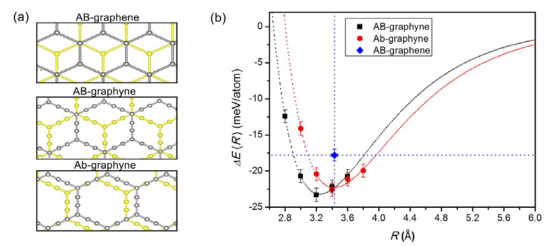

Figure 8.3 (a) Stacking configurations of AB-stacked bilayer graphene and two stable modes (AB and Ab) of a bilayer α-graphyne. The yellow and gray structures represent the lower and upper layer of a bilayer, respectively. (b) Diffusion Monte Carlo (DMC) interlayer binding energies of AB- and Ab-stacked bilayer α-graphynes as functions of an interlayer distance. The blue diamond symbol represents a DMC interlayer binding energy for an AB-bilayer graphene at the equilibrium interlayer distance.

(Image courtesy of Shin, H., et al., *J. Chem. Theory Comput.* **13**, 5639 [2017].)

The Impact

This work is a predictive study of the geometry and energy properties of graphyne using the highly accurate quantum Monte Carlo (QMC) method as implemented in the QMCPACK code. When comparing the bilayer binding to that found in graphene (bound only by van der Waals [vdW] interactions), covalent interactions are found to enhance the binding of the bilayers in graphyne. These QMC results were used to assess the accuracy of vdW-corrected density functional theory (DFT) functionals with respect to predicting the geometries and properties of this class of materials.

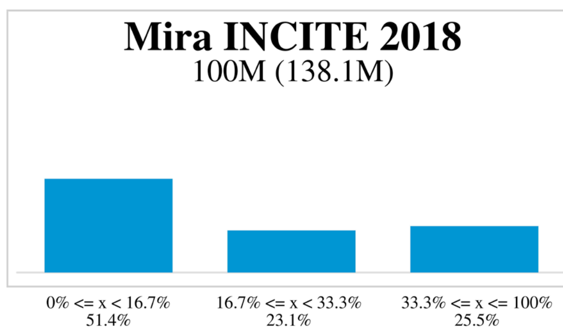

Summary

α-Graphyne is a two-dimensional (2D) sheet of sp−sp2 hybridized carbon atoms in a honeycomb lattice (Figure 8.3). While the geometrical structure is similar to that of graphene, the hybridized triple bonds give rise to an electronic structure that is different from that of graphene. The band structure of α-Graphyne is expected to have a valence band and a conduction band taking the shape of the upper and lower halves of a conical surface, also known as a Dirac cone, and meeting at a specific point known as the Dirac point. Materials with Dirac cones in their band structures present ultra-high carrier mobility and can lead to very high-frequency semiconductors. While not yet synthetized, its existence was predicted in the early 1960s. Since then, many *ab initio* studies have tried to describe the properties and the most stable geometry of graphyne. The widely used DFT does not take into account dispersion forces, which are key to describing vdW forces and therefore need to be corrected with empirical parameters when used on vdW-dominated systems. These corrections are usually fitted to experimental measurements; but in the absence of such measurements, more accurate theories are needed to provide reliable references. Using QMC, this study was able to accurately describe this carbon allotrope and assess the stability of its various stacking modes while comparing it to the well-known graphene allotrope. These high-quality results enabled the researchers to assess the accuracy of recent vdW-corrected DFT functionals, showing their significant bias toward underestimated interlayer binding energy and overestimated equilibrium distances. Finally, it was shown that because of the very small energy difference between graphene and graphyne, temperature effects must play a significant role in the relative stability of the two allotropes, yielding a valuable insight to experimentalists seeking to synthesize this compound. ALCF contributions include ALCF staff and postdocs contributing to the development of QMCPACK as part of the two Early Science Program projects and in support of the INCITE program. In collaboration with other DOE laboratories and with support from external funding sources (DOE-Basic Energy Sciences [BES], ECP, Intel-Parallel Computing Centers [PCCs]), staff and postdocs contributed to new science methods and numerous software engineering enhancements.

Contact

Paul R. C. Kent
Oak Ridge National Laboratory
kentpr@ornl.gov

Publications

Shin, H., et al. "Nature of Interlayer Binding and Stacking of sp–sp2 Hybridized Carbon Layers: A Quantum Monte Carlo Study." *J. Chem. Theory Comput*. **13**, 5639 (2017).

Kim, J., et al. "QMCPACK: An Open Source Ab Initio Quantum Monte Carlo Package for the Electronic Structure of Atoms, Molecules, and Solids." *J. Phys. Condens. Matter* **30** (19).

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR

July 2018

8.2.4 Global Radiation Magnetohydrodynamics (MHD) Simulations of Massive Star Envelopes

The Science

Massive stars live relatively short but dramatic lives. Stars like our sun can shine steadily for billions of years, whereas massive stars can live for only millions of years or less, and their luminosity and color can vary significantly on timescales of days to years. Simulations that try to approximate the outer layers of these massive stars as spherically symmetric are not able to model this behavior well enough to be predictive. This constraint limits our understanding of the life cycles of massive stars, which in their deaths are the progenitors of some of the most interesting and energetic phenomena in the universe, including supernovae, gamma-ray bursts, neutron stars, and black holes.

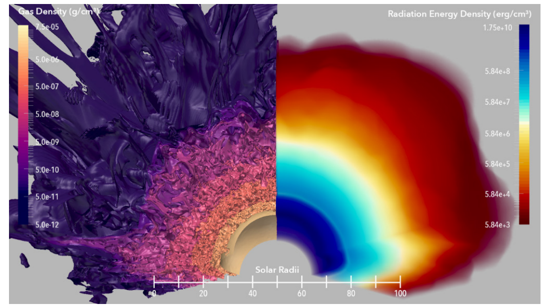

Figure 8.4 Outer layers of a massive variable star from a simulation run on Mira. Left: a slice through the 3D density distribution of matter, clearly showing the effects of turbulence. Right: a slice through the 3D distribution of energy in photons.

(Image courtesy of Joseph A. Insley, Argonne National Laboratory; and Yan-Fei Jiang, University of California, Santa Barbara.)

The Impact

By harnessing the power of 8,192 nodes of Mira, the research team found that fully 3D simulations of the outer layers of massive stars were dramatically improved in their ability to model some observed behaviors. An initial simulation campaign was able to reproduce the variability patterns of luminous blue variable (LBV) stars on timescales of weeks to months and also identified the critical physical mechanisms involved. The researchers also found a correlation between temperature and variability patterns on the timescale of days that could be verified with future high-cadence observations of LBVs. This work paves the way to using similar simulation methods to understand the behaviors of other kinds of massive stars.

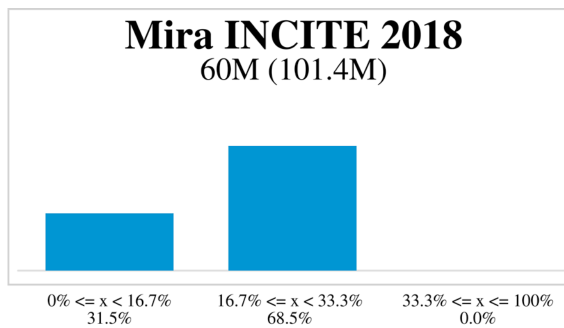

Summary

This project set out to overcome the limitations of one-dimensional simulations of massive stars in modeling these stars' variability and outburst activity, which, in turn, prevents those simulations from having predictive power. The group used the 1D MESA (Models for Experiments in Stellar Astrophysics) code to generate initial distributions of matter and radiation in the outer layers of massive stars. Those distributions served as the initial conditions for fully 3D simulations run using Athena++, a grid-based code that models compressible magnetohydrodynamics (MHD) and time-dependent radiative transfer. One of the simulation configurations naturally reproduced the luminosity, effective temperature, and mass loss rate of LBV stars during outburst activity. The team identified two key physical processes that can only be modeled in 3D simulations: (1) pressure support in the extended envelope of the star from supersonic turbulence, and (2) energy transport through turbulent opacity peak regions via radiation and convection (Figure 8.4). By comparing the simulation that accurately reproduced observed LBV outburst properties with several other initial configurations that did not, the authors were able to highlight the important role that helium opacity plays in triggering the outbursts and setting the observed effective temperature. The simulations also predict a correlation between the effective temperature of the LBV and the amplitude of irregular variability on the timescale of days. Such short time-scale variability has recently been observed, and further high-cadence observations could potentially verify the predicted correlation. ALCF staff contributed the following: (1) helped with a critical redesign of the file reading and writing infrastructure of the Athena++ code used for this project and achieved a 100× improvement in file I/O performance at capability scale on Mira, (2) helped generate movies and still images of production simulations used for presentations and publications, (3) helped bundle simulations together for improved queue throughput, and (4) consulted on hybrid parallelism as a programming technique to address memory scaling issues.

Contact

Lars Bildsten
Kavli Institute for Theoretical Physics
University of California, Santa Barbara
bildsten@kitp.ucsb.edu

Publication

Jiang, Y.-F., et al. "Outbursts of Luminous Blue Variable Stars from Variations in the Helium Opacity." *Nature* **561**, 498–501 (2018). [DOI:10.1038/s41586-018-0525-0]

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR

October 2018

8.2.5 Understanding Optical Properties of Layered Perovskites with Extreme-Scale Computing

The Science

Hybrid organic-inorganic layered perovskites (HOIPs) are semiconductor materials with opto-electronic properties that can either produce or capture light and are excellent candidates for low-cost solar cells given their enhanced charge transport properties. Laboratory experiments of HOIPs suggested that perovskites could be designed to have tailored properties; however, the nature of combining potential materials makes this a daunting task. Although researchers have used empirical models to explain the fundamental properties of HOIPs, they traditionally needed more complex experimental results or theoretical studies for validation. This team has shown that the electronic properties of HOIPs can be quantitatively predicted and modulated (by separately modifying the organic and inorganic portions) and are well described by a quantum-well model.

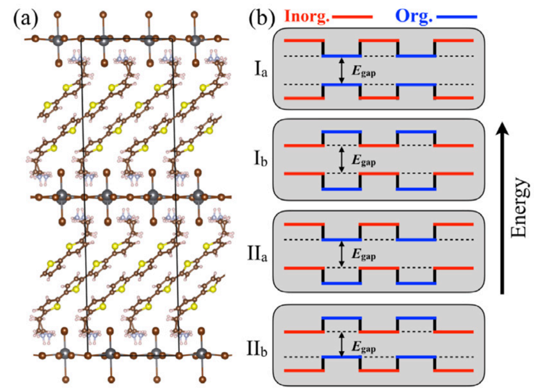

Figure 8.5 (a) The supercell structure of a layered hybrid perovskite. Planar layers represent inorganic structures, whereas organic chains are ordered in diagonal arrangements. (b) Possible energy-level schemes for an alternating organic-inorganic perovskite are shown with the observed band gap indicated by the arrows and dashed lines.

(Image courtesy of Volker Blum, Duke University; Figure 1 in *Phys. Rev. Lett.* **121**, 146401 [2018].)

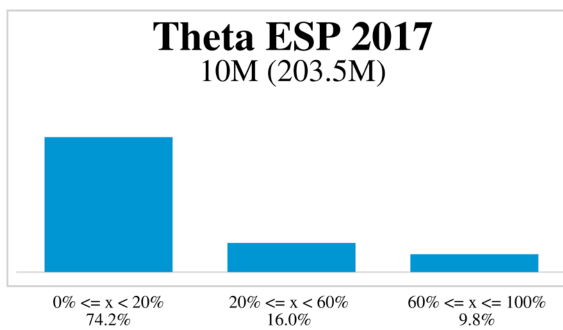

The Impact

There is a technology rush for materials to produce electronic components that efficiently generate electricity from sunlight. HOIPs are extremely good at performing this task and, furthermore, could be combined with powerful light-emitting devices to create low-cost devices that emit light. Enabled by Theta at the ALCF, it was possible for the first time to calculate — with an all-electron quantum mechanics method — the electronic structure of complex HOIP systems of sufficient sizes to corroborate experimental results and show that quantum-well models can be useful tools to design HOIPs. These findings open new paths for modeling and tailoring the properties of HOIPs to aid in the design and targeted synthesis of materials in the laboratory.

Summary

Empirical models suggested that researchers could layer hybrid perovskites, and the interfaces created would give rise to an electronic structure that formed quantum wells (localized regions of low potential energy). By design, these quantum wells could be staggered to produce an array of electron carriers in the presence of light by which charge could easily be transported. To confirm this important finding, a series of systematic and highly accurate calculations were needed, in particular to separately adjust the sizes of the organic and inorganic layers to tune the confinement and alignment of the "quantum-wells-like" model. A team led by materials scientist Volker Blum at Duke University carried out extensive all-electron, periodic-boundary-condition calculations of hybrid perovskites using hybrid, spin-orbital, and dispersion-corrected DFT methods. The calculations, utilizing experimental structures of $AE4TPbX_4$ (X=Cl, Br, I; see Figure 8.5) and a paradigmatic example of a 2D HOIP as starting points, provided insight into why the X=Cl analog had a strong emission, whereas the X=Br, I systems were substantially quenched. These computationally expensive calculations were possible thanks to the high scalability of the FHI-AIMS code (originally developed by Blum), the use of an efficient eigenvalue solver, and the unique architectural hardware features of Theta, which was accessed through the ALCF's Theta ESP. The present computational work is an important first step in solving longstanding questions to support laboratory observations using large–scale, state-of-the-art electronic structure modeling on large-scale DOE computational resources. The ALCF helped to port the FHI-AIMS code and optimize some functions, which enabled Theta to perform large-scale linear algebra computations.

Contact

Volker Blum
Duke University
volker.blum@duke.edu

Publication

Liu, C., *et al.*, "Tunable Semiconductors: Control Over Carrier States and Excitations in Layered Hybrid Organic-Inorganic Perovskites." *Phys. Rev. Lett*. **121**, 146401 (2018).
[DOI: 10.1103/PhysRevLett.121.146401]

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR

November 2018

8.2.6 Atomically Thin Materials with Engineered Strain

The Science

Merging multiple units of different crystal structures could lead to new materials with outstanding properties. This project showed that it is possible to fuse materials, atom-by-atom, to produce atomic-scale fabrics with nearly flawless interfaces. These nanofabrics have the potential for demonstrating extraordinary properties for use in diverse electronic applications, such as in light-emitting diodes.

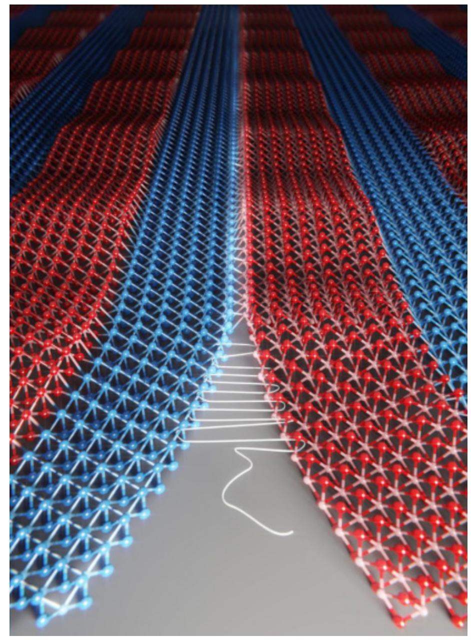

Figure 8.6 Schematic representation of two distinct, 1-atom-thick materials (red and blue) being "sewn" together to create a new material with improved properties.

(Image courtesy of Jiwoong Park, University of Chicago.)

The Impact

Future generations of portable electronics, including mobile sensors and wearable technology, will need to be more efficient in power consumption and reduced in size. This study addresses the problem of nanofabrication by manipulating at the atomic-scale transition metal layers, effectively sewing together 1-atom-thick fabrics to produce the building blocks for micro-electronic components.

Summary

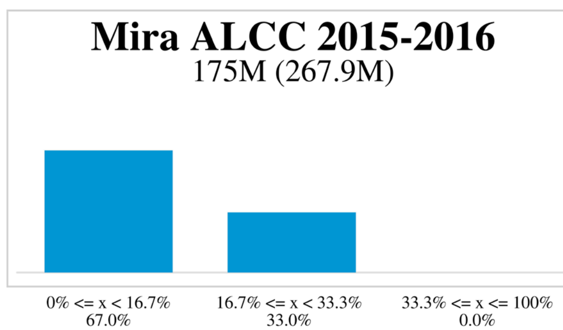

This project, led by Jiwoong Park (University of Chicago), showed a new path to bringing atom-by-atom crystal structures together to create efficient 1-atom-thick conductors (Figure 8.6) with great potential to produce portable electronic devices that are lighter with improved energy consumption as compared with current technologies. The experimental-computational team reports that it is possible to create precise single-atom–thick, coherent 2D structures and super lattices to produce nanoscale building blocks for electronic devices, allowing fine-tuning of macroscopic properties (e.g., optical and conductive). The computational component of this project was led by Robert A. DiStasio, Jr., and performed by his group at Cornell University. The simulations required were a series of coarse-grained models to define the atoms' arrangements and electronic structure calculations to estimate long-range interactions on the deposition of crystal monolayers on surfaces. The team required large-scale computational resources because of the large number of particles needed and the increased computational expense of the DFT-based vdW model. Because of the complexity of running multiple simulations concurrently with dependencies, ALCF staff helped create ensemble job submission scripts to pack multiple tasks into capability jobs using job dependencies to improve throughput in the queues.

Contact

Robert A. DiStasio, Jr.
Cornell University
distasio@cornell.edu

Publication

Xie, S., *et al*. "Coherent, Atomically Thin Transition-Metal Dichalcogenide Superlattices with Engineered Strain." *Science* **359**, 1131–1136 (2018). [DOI:10.1126/science.aao5360]

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR

September 2018

8.2.7 Extreme–Scale, Unstructured Adaptive CFD: From Multiphase Flow to Aerodynamic Flow Control

The Science

Understanding the flow of fluids through solid bodies is a long-standing challenge. Many of the great minds of science have built up a hierarchy of fluid models to understand the complex fluid phenomena that very often involve turbulent flows. At extreme fluid flow situations, turbulent flows develop extremely broad ranges of length and timescales. This disparity motivates the use of discretization methods to solve the Navier-Stokes equations. The combination of features poses a challenge to optimal implementation on massively parallel platforms. However, the more complex algorithms can provide great reductions in computational cost relative to simpler methods.

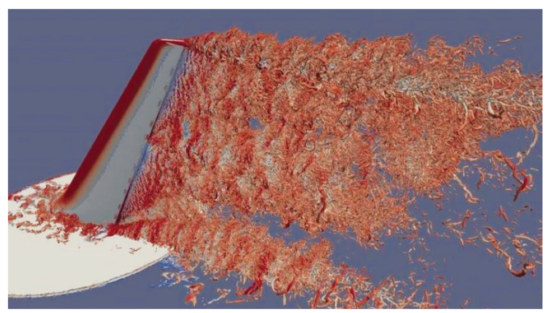

Figure 8.7 Visualization of vortical structures from synthetic jet actuators that affect flow control on a model aircraft's vertical tail.

(Image courtesy of Kenneth Jansen, University of Colorado, Boulder.)

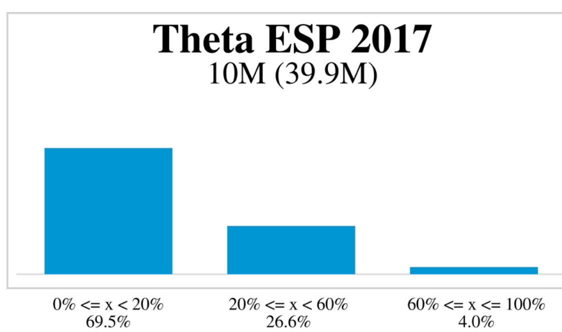

The Impact

The energy impact of the predictive capability that has been developed in this project can best be related to its potential to reduce the size (and weight) of the vertical tail and rudder, and also reduce their significant drag contribution during cruise conditions. The joint experimental/computational studies suggest that active flow control can achieve this size reduction at experiment scale.

Summary

This Theta ESP project was aimed primarily at realizing the goal of improved aerodynamic design and aircraft performance when lift is maximized and drag and the weight of the aircraft components are minimized, all of which contribute to a reduction in the aircraft's fuel consumption. A recent study at Boeing estimated that a 777-class airplane could reduce its fuel consumption by 0.75–1.0% on a 3,000-nautical-mile trip if its vertical tail size could be reduced by 25%. By replacing the control surface on the vertical tail with synthetic jet actuators for active flow control, it has been shown that large-scale flow changes (e.g., reattachment of separated flow or virtual aerodynamic shaping of lifting surfaces) can be produced from microscale input (e.g., a 0.1-W piezoelectric disk resonating in a cavity), thereby resulting in a substantial reduction in the weight of the control surfaces.

A scalable framework for simulation and data analysis has been developed for simulating the flow over the vertical tail surface of an aircraft with 12 synthetic jets using highly resolved detached eddy simulation (DES) models that were implemented in the massively parallel PHASTA finite element solver. The simulations have shown excellent agreement not only with respect to integrated quantities like total force, but also phase-averaged flow structures that issue from the synthetic jets (see Figure 8.7) yielding, for the first time, clear insights into the fundamental mechanisms of flow control. The research team is now in the process of extending these experiment-scale simulations to an eight-times-higher Reynolds number, which will bring the simulations substantially closer to the Reynolds number at *actual* flight conditions. These simulations can help us understand how the flow control structures and the jets that create them must be adjusted for the Reynolds number. This capability now sets the stage for the true flight-scale simulations that Aurora will make possible. ALCF staff collaborated with the research team on improving the performance of PHASTA, which was coupled to anisotropic adaptive meshing procedures.

Contact

Kenneth E. Jansen
University of Colorado, Boulder
kenneth.jansen@colorado.edu

Publications

Jansen, K., *et al*., "Extreme Scale Unstructured Adaptive CFD: From Multiphase Flow to Aerodynamic Flow Control." Theta ESP Report, ANL/ALCF/ESP-17/1, Argonne Leadership Computing Facility (2017).

Jansen, K.E., *et al.*, "Interaction of a Synthetic Jet Actuator on Separated flow over a Vertical Tail." 35th AIAA Applied Aerodynamics Conference, AIAA AVIATION Forum, (AIAA 2017-3243) (2017).

Fang, J., *et al*., "Interface Tracking Simulations of Bubbly Flows in PWR Relevant Geometries." *Nuc. Eng. Design* **312**, 205–213 (2017). [DOI: 10.1016/j.nucengdes.2016.07.002]

Fang, J., *et al*., "Interface Tracking Investigation of Geometric Effects on the Bubbly Flow in PWR Subchannels." *Nucl. Sci. Eng.* **193**, 46–62 (2019). [DOI:10.1080/00295639.2018.1499280]

Highlight Categories
Performer/Facility: ASCR-ALCF

Date Submitted to ASCR
April 2018

8.3 Stakeholder Engagement

To help ensure that the ALCF delivers on its mission of delivering breakthrough science, staff outreach needs to closely engage with domain science and keep 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 machine architectures in the near term and in the future.

In 2018, 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 JLSE and for ALCF projects relevant to facility and application readiness. (JLSE activities are detailed in Section 4.2.1 of this report.)

ALCF staff support a wide range of computer science and domain science projects, and work in close collaboration with the project teams to advance their use of both production resources and future resources alike. Additionally, staff members participate in community and domain activities, especially at the annual SC conference, where ALCF staff members lead numerous sessions including technical talks, Birds of a Feather, workshops, and tutorials. Staff also visit universities and participate in town hall conversations 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. Notably, in 2018, ALCF staff organized and ran LLVM-HPC2018: The Fifth Workshop on the LLVM Compiler Infrastructure in HPC, at SC18.

8.3.1 Outreach

Engagement in Standards and Community Groups

The ALCF participates in a number of HPC standards and community groups in order to promote ALCF interests, educate the community about ALCF resources, and increase collaboration with ALCF staff. These include the following: HCP User Forum; Cray User Group; Intel Xeon Phi User Group (IXPUG); C++ Standards Committee; and MPI Forum.

Industry Lunch at SC Conference

The ALCF's industry partnerships and outreach manager organizes an annual lunch at SC for ALCF senior leadership and select industry users of HPC to discuss ways to effectively work together. New attendees to the 2018 lunch included Argonne's MCS division director and the head of Argonne's Laboratory Center for Research Computing. Five of the eight industry attendees are current ALCF users. Topics of discussion included training, software availability, and emerging technologies such as quantum and deep learning.

Intel Xeon Phi User Group

The Intel Extreme Performance User Group (IXPUG) is an independent user group whose members come from sites with major Intel Xeon Phi installations. In 2018, the ALCF's industry partnerships and outreach manager, in his capacity as current IXPUG president, led the group's transition from a Xeon-specific focus to one that addresses the full range of HPC offerings in Intel's portfolio. IXPUG held eight significant meetings in 2018, including a spring conference in Bologna, Italy, and a fall conference in Hillsboro, Oregon. In addition, IXPUG held workshops and birds-of-a-feather sessions at the International Supercomputing Conference (ISC) and at SC. ALCF staff also organized and hosted the 2018 IXPUG In Situ Visualization Hackathon in July.

Lustre Users Group 2018

The ALCF hosted the annual Lustre Users Group (LUG) 2018 meeting. LUG brings together Lustre users from industry, academia, and government, as well as storage vendors to meet and discuss Lustre. The event attracted 130 attendees from both inside and outside the United States by offering the opportunity to learn from each other and trade best practices.

Collaboration Meeting with Northwestern University

On May 21, 2018, ALCF staff attended the Northwestern-Argonne Institute of Science and Engineering (NAISE) open house to discuss ideas for new research opportunities, including routes to successful proposals for obtaining computational time at the ALCF, and to bring respective researchers together to build new collaborations. NAISE is an effort to advance Northwestern University's (NU) and Argonne's scientific goals while expanding education and research opportunities for undergraduate and graduate students.

Northwestern-Argonne Workshop on Computational Neuroscience

During 2018, Argonne was interested in developing stronger links between the neuroscience community at NU and staff members from the ALCF and Computational Science (CPS) division in collaboration with NAISE and the Northwestern University Interdepartmental Neuroscience Program. Argonne hosted a one-day workshop at the lab to bring their respective researchers together to build new collaborations. The Northwestern-Argonne Workshop on Computational Neuroscience was held on August 31, 2018, at Argonne's Theory and Computing Science Building.

Quantum Computing Training

ALCF staff members organized a national Quantum Computing Workshop, held at Argonne, July 25–27, 2018. The workshop goals were to inform attendees about the quantum computing programs at participating institutions and highlight promising research opportunities; to serve as an incubator for future collaborations; to teach attendees about available tools for simulating quantum computers; and to present state-of-the-art research in quantum computing. The workshop brought together more than 150 representatives from across industry to discuss existing programs and projects and to examine promising research opportunities. Speakers included representatives from The University of Chicago, Intel, Google, Rigetti, the DOE national laboratories, and more.

Resources for Performance Portability

In 2017, an ALCF senior staff member 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. This site has been maintained, with new content added throughout 2018.

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 2018, 27 students worked on projects ranging from system administration and data analytics to computational science and performance engineering.

Fluid Dynamics Workshop

An ALCF staff member organized the Workshop on High Reynolds Number Flow Simulations on Exascale Platforms, held at Argonne on September 17–19, 2018. The workshop goals were to advance the state of the art of turbulence modeling and simulation of high Reynolds number flows, and to build a community of researchers who use and develop a set of open-source simulation codes, software tools for uncertainty quantification, *in situ* flow visualization, machine learning, and feature extraction methods on Aurora and beyond. Approximately 48 people attended, which included Argonne staff, U.S. and European university faculty members, agency program managers, government employees (DOE, U.S. Department of Defense, National Aeronautics and Space Administration), and a researcher from a U.S. aerospace company.

8.3.2 Summary of Engagements with the Exascale Computing Project

Argonne is a core laboratory of the ECP, and every member of the ALCF leadership team is engaged in the ECP project at some level. Susan Coghlan and David Martin are on the ECP leadership team: Coghlan is deputy director of Hardware and Integration (HI) and Martin is co-executive director of the Industry Council. (Additionally, Argonne's Andrew Siegel is director of Applications Development.) Other leadership team members participate in the various working groups and projects, including Katherine Riley (Application Development), Mark Fahey (Facilities), Jini Ramprakash (Facilities), and Scott Parker (Software Technology). ALCF Division Director Michael E. Papka regularly participates in teleconferences with the ECP project director and other facility directors. In addition, numerous other ALCF staff members have roles in the projects and working groups listed above, as well as PathForward and Continuous Integration, or CI, efforts.

Last year, 19 ALCF staff attended the ECP Annual Meeting held February 5–9, 2018, in Knoxville, Tennessee, to participate in technical conversations, project discussions, and facility-specific breakouts. In addition, the ALCF participated in several planning meetings with ECP and the other computing facilities (NERSC, OLCF) to develop and deploy the ECP/Facilities engagement plan, and worked with the ECP training lead to promote ECP training activities to ALCF users.

The ALCF ECP/HI Applications Integration effort made great strides in 2018 by hiring multiple performance engineers and training staff to support application readiness. Finally, in coordination with ECP HI, the ALCF hired staff to work on CI for ECP. CI is the software engineering practice of automatically testing and integrating changes to a project's code before they are released to end users — a process that involves merging all developers' working copies to a shared mainline several times a day.

With more user communities at ALCF adopting CI into their workflows, the ALCF proposed a two-prong approach, a tactical solution and a strategic solution, to meet the needs of a wide range of unique user requirements and to connect to distributed software repositories. The second prong, the strategic solution, is the deployment of the ECP CI solution based on GitLab currently under development and earmarked for Q2 2019. The ALCF is participating in the ECP CI working group during its development.

Use of Jenkins Automation Server for CI

The use of GitLab should provide an easier CI experience to the end user. A stock GitLab installation allows an organization to host code repositories, build code using GitLab runners, and manage the security polices of each project (Figure 8.8). ECP CI is developing several enhancements to GitLab for HPC environments, including the ability for GitLab runners to perform build tasks and to interact with HPC schedulers, and seamless integration across various DOE sites so that projects can perform friction-free testing against different HPC architectures.

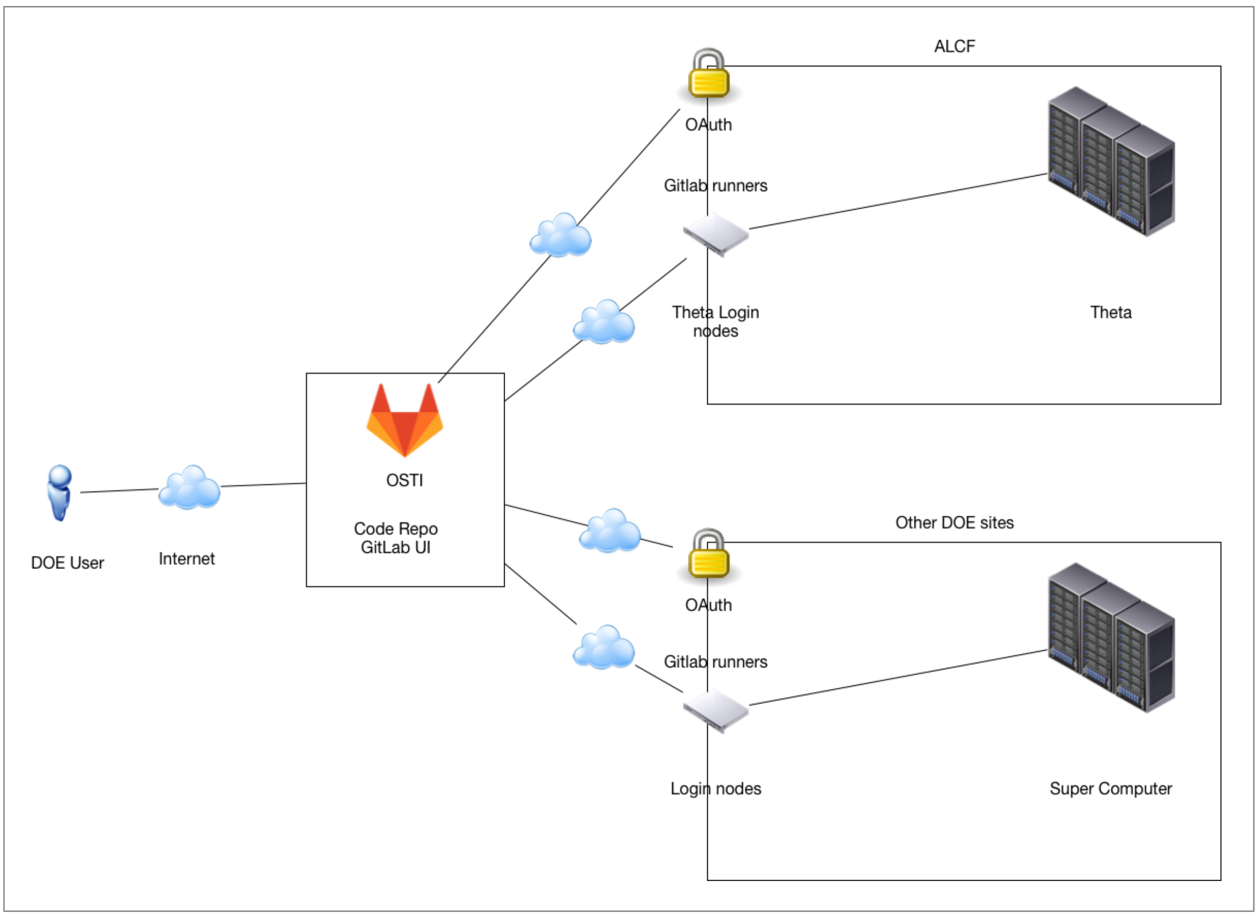

Figure 8.8 High-level Process Flow for the Use of Continuous Integration at the ALCF and Other Sites

The current Jenkins solution and the long-term GitLab solution provide our users with the ability to build and test their code at a given interval or event. This testing enables teams to obtain quicker feedback when their code is negatively affected by their changes.

Communication between the ALCF and the ECP Resource Allocation Council

In 2018, the ECP ALCC allocation completed and the compute facilities and ECP switched to the RAC (discussed in Section 3.1.2) to support ECP computing needs. This group, composed of representatives of the facilities and the ECP, meets monthly to review project progress and to assess new project needs.

To help automate how the RAC consumes this data, the ALCF also sends a .CSV file (comma-separated values file) report to the ECP mailing list each week (Figure 8.9). While the web version is visually appealing (Figure 8.10), the .CSV format can be merged into any other workflow to process.

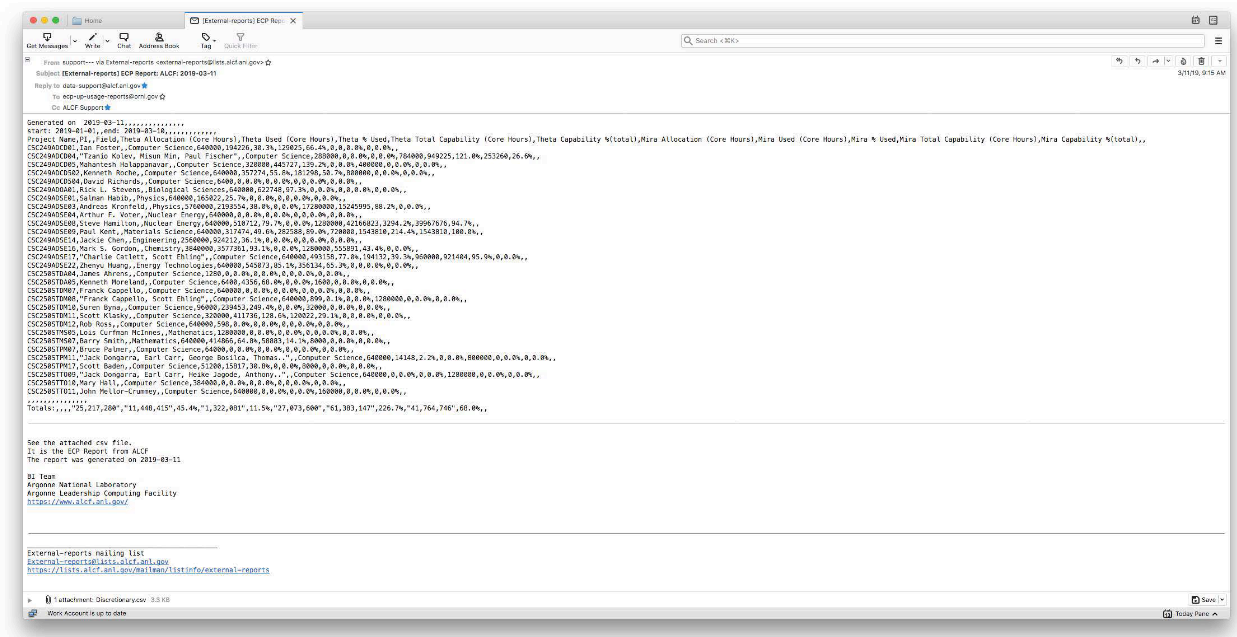

Figure 8.9 Exascale Computing Project Weekly E-mail Screen Shot

ECP Discretionary 01/01/2019 to 03/31/2019 (subject to revision)

Generated on 2019-02-13 Export to Excel Export to CSV

			THETA					MIRA				
Project Name	**PI**	**Field**	**Allocation (Core Hours)**	**Hours Used (Core Hours)**	**% Used**	**Total Capability (Core Hours)**	**Capability %(total)**	**Allocation (Core Hours)**	**Hours Used (Core Hours)**	**% Used**	**Total Capability (Core Hours)**	**Capability %(total)**
CSC249ADCD01	Ian Foster	Computer Science	640,000	59,052	9.2%	47,623	80.6%	0	0	0.0%	0	0.0%
CSC249ADCD04	Tzanio Kolev, Misun Min, Paul Fischer	Computer Science	288,000	0	0.0%	0	0.0%	784,000	949,225	121.0%	253,260	26.6%
CSC249ADCD05	Mahantesh Halappanavar	Computer Science	320,000	0	0.0%	0	0.0%	400,000	0	0.0%	0	0.0%
CSC249ADCD502	Kenneth Roche	Computer Science	640,000	656	0.1%	0	0.0%	800,000	0	0.0%	0	0.0%
CSC249ADCD504	David Richards	Computer Science	6,400	0	0.0%	0	0.0%	0	0	0.0%	0	0.0%
CSC249ADOA01	Rick L. Stevens	Biological Sciences	640,000	8,820	1.3%	0	0.0%	0	0	0.0%	0	0.0%
CSC249ADSE01	Salman Habib	Physics	640,000	165,022	25.7%	0	0.0%	0	0	0.0%	0	0.0%
CSC249ADSE03	Andreas Kronfeld	Physics	640,000	771,042	120.4%	0	0.0%	1,280,000	2,565,466	200.4%	0	0.0%
CSC249ADSE04	Arthur F. Voter	Nuclear Energy	640,000	0	0.0%	0	0.0%	0	0	0.0%	0	0.0%
CSC249ADSE08	Steve Hamilton	Nuclear Energy	640,000	0	0.0%	0	0.0%	1,280,000	1,018,138	79.5%	0	0.0%
CSC249ADSE09	Paul Kent	Materials Science	640,000	14,150	2.2%	13,380	94.5%	720,000	0	0.0%	0	0.0%
CSC249ADSE14	Jackie Chen	Engineering	640,000	719,018	112.3%	0	0.0%	0	0	0.0%	0	0.0%
CSC249ADSE16	Mark S. Gordon	Chemistry	3,840,000	2,251,450	58.6%	0	0.0%	1,280,000	493,823	38.5%	0	0.0%
CSC249ADSE17	Charlie Catlett, Scott Ehling	Computer Science	640,000	279,493	43.6%	0	0.0%	960,000	8,392	0.8%	0	0.0%
CSC249ADSE22	Zhenyu Huang	Energy Technologies	640,000	65,573	10.2%	0	0.0%	0	0	0.0%	0	0.0%
CSC250STDA04	James Ahrens	Computer Science	1,280	0	0.0%	0	0.0%	0	0	0.0%	0	0.0%
CSC250STDA05	Kenneth Moreland	Computer Science	6,400	4,356	68.0%	0	0.0%	1,600	0	0.0%	0	0.0%
CSC250STDM07	Franck Cappello	Computer Science	640,000	0	0.0%	0	0.0%	0	0	0.0%	0	0.0%
CSC250STDM08	Franck Cappello	Computer Science	640,000	0	0.0%	0	0.0%	1,280,000	0	0.0%	0	0.0%
CSC250STDM10	Suren Byna	Computer Science	96,000	239,453	249.4%	0	0.0%	32,000	0	0.0%	0	0.0%
CSC250STDM11	Scott Klasky	Computer Science	320,000	410,438	128.2%	120,022	29.2%	0	0	0.0%	0	0.0%
CSC250STDM12	Rob Ross	Computer Science	640,000	522	0.0%	0	0.0%	0	0	0.0%	0	0.0%
CSC250STMS05	Lois Curfman McInnes	Mathematics	1,280,000	0	0.0%	0	0.0%	0	0	0.0%	0	0.0%
CSC250STMS07	Barry Smith	Mathematics	640,000	19,286	3.0%	0	0.0%	8,000	0	0.0%	0	0.0%
CSC250STPM07	Bruce Palmer	Computer Science	64,000	0	0.0%	0	0.0%	0	0	0.0%	0	0.0%
CSC250STPM11	Jack Dongarra, Earl Carr, George Bosilca, Thomas..	Computer Science	640,000	14,148	2.2%	0	0.0%	800,000	0	0.0%	0	0.0%
CSC250STPM17	Scott Baden	Computer Science	51,200	0	0.0%	0	0.0%	8,000	0	0.0%	0	0.0%
CSC250STTO09	Jack Dongarra, Earl Carr, Heike Jagode, Anthony..	Computer Science	640,000	0	0.0%	0	0.0%	1,280,000	0	0.0%	0	0.0%
CSC250STTO10	Mary Hall	Computer Science	384,000	0	0.0%	0	0.0%	0	0	0.0%	0	0.0%
CSC250STTO11	John Mellor-Crummey	Computer Science	640,000	0	0.0%	0	0.0%	160,000	0	0.0%	0	0.0%
		Totals:	**18,177,280**	**5,022,477**	**27.6%**	**181,025**	**3.6%**	**11,073,600**	**5,035,044**	**45.5%**	**253,260**	**5.0%**

·Theta Capability: Jobs that used at least 20% of the machine

·Mira Capability: Jobs that used at least 16.7% of the machine

Figure 8.10 Exascale Computing Project Website Screen Shot

Conclusion

The ALCF continues to enable scientific achievements, consistent with DOE's strategic goals for obtaining scientific breakthroughs and advancing the foundations of science through projects carried out on facility machines. Researchers participating in projects using ALCF resources published 280 papers in CY 2018. ALCF projects have had success in a variety of fields and using many different computational approaches. Our users 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.

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:

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

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	799 / 51,136
Yes	No	20.0% <= x < 60.0%	800 / 51,200	2,399 / 153,536
Yes	Yes	60.0% <= x	2,400 / 153,600	See: A.7 Theta Nodes

Capability also refers to a calculation. The capability calculation is the percentage 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: Innovative and Novel Computational Impact on Theory and Experiment (INCITE) capability is the percentage of core-hours of INCITE jobs with the capability attribute versus the total core-hours of all INCITE jobs for a time period.

***Formula*:**

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

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

A.7 Theta Nodes

The number of reportable nodes on Theta is fewer than the total number of 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 – Argonne Leadership Computing Facility's Director's Discretionary Projects: Mira

January 1, 2018 – December 31, 2018
Director's Discretionary (DD) Projects on Mira

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
AEM_Hydroxide_NYU	Mark E Tuckerman	New York University	First-principles Discovery of Design Rules for Anion Exchange Membranes with High Hydroxide Conductivity	Chemistry	7,000,000
ALCF_Getting_Started	Ray Loy	Argonne National Laboratory	Argonne Leadership Computing Facility (ALCF) Getting Started	Training	10,000
AlgLA	Oded Schwartz	The Hebrew University	Algorithmic Linear Algebra	Computer Science	1,250,000
Allinea	Ray Loy	Argonne National Laboratory	Improved Debugging Memory Usage for Blue Gene/Q (BG/Q)	Internal	1,000,000
AMASE	Rajkumar Kettimuthu	Argonne National Laboratory	Architecture and Management for Autonomous Science Ecosystems	Computer Science	100,000
ATPESC18_Instructors	Marta Garcia Martinez	Argonne National Laboratory	Argonne Training Program on Extreme Scale Computing for ALL Instructors	Training	100,000
ATPESC2018	Marta Garcia Martinez	Argonne National Laboratory	Argonne Training Program on Extreme-Scale Computing 2018	Computer Science	10,000,000
AstroGK_turb	Yohei Kawazura	University of Oxford	Energy Flow and Heating in Astrophysical Plasma Turbulence	Physics	1,000,000
athena_performance	Brian OShea	Michigan State University	Scaling and Performance Enhancement of an Astrophysical Plasma Code	Physics	1,000,000
aurora_app	William Scullin, Paul Coffman	Argonne National Laboratory	Aurora Application Enablement	Computer Science	6,000,000
AutomatedBench	Kevin Harms	Argonne National Laboratory	Automated Benchmarking	Computer Science	2,500,000
bloodflow_dd	Jifu Tan	Northern Illinois University	Microfluidic Design and Optimization for Cell Separation	Engineering	2,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
Catalyst	Katherine Riley	Argonne National Laboratory	Catalyst	Internal	20,000,000
CFD-TPM	Prasad Vegendla	Argonne National Laboratory	Validation of Two-phase Flow Models and Critical Heat Flux Prediction for the Highly Scaleable Computational Fluid Dynamics (CFD) Code NEK-2P	Nuclear Energy	5,000,000
Carbon_composites	Hendrik Heinz	University of Colorado-Boulder	Rational Design of Ultrastrong Composites	Materials Science	2,500,000
climate_severe	Victor Gensini	Northern Illinois University (NIU)	Anticipating Severe Weather Events via Dynamical Downscaling	Earth Science	3,000,000
CCV_HPC4Mfg	Muhsin M Ameen	Argonne National Laboratory	CFD Study of Impact of Part-to-Part Variations on Spark-Ignition Engine Charge Formation	Engineering	5,000,000
CharmRTS	Laxmikant V. Kale	University of Illinois at Urbana-Champaign	Charm++ and Its Applications	Computer Science	2,808,887
CobaltDevel	Paul Rich	Argonne National Laboratory	Cobalt Development	Internal	10,000,000
CNTmetallization	Iman Salehinia, Michael E. Papka	Northern Illinois University	Metallization of Carbon Nanotubes (CNTs) for Thermal and Structural Applications	Materials Science	2,000,000
CompSpectro_DD	Giulia Galli	The University of Chicago	Electronic and Transport Properties of Complex Interfaces	Materials Science	40,000,000
CORALDev	Scott Parker	Argonne National Laboratory	CORAL Development and Testing	Internal	6,000,000
Comp_Perf_Workshop	Ray Loy	Argonne National Laboratory	Comp_Perf_Workshop	Training	5,000,000
critical_perf	Scott Parker, Ray Loy	Argonne National Laboratory	Critical Debugging Project	Internal	20,000,000
CrystalsADSP	Alexandre Tkatchenko, Alvaro Vazquez Mayagoitia	University of Luxembourg	Constructing and Navigating Polymorphic Landscapes of Molecular Crystals	Materials Science	85,000,000
CSC249ADCD04	Tzanio Kolev, Misun Min, Paul Fischer	Lawrence Livermore National Laboratory (LLNL)	1.2.5.3.04 Center for Efficient Exascale Discretizations (CEED)	Computer Science	2,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CosmicLaser_DD	Petros Tzeferacos	The University of Chicago	Simulations of Laser Experiments to Study the Origin of Cosmic Magnetic Fields	Physics	5,000,000
CSC249ADCD02	Timothy Germann	Los Alamos National Laboratory	2.2.6.04 ADCD02-COPA: Co-Design Center for Particle Applications	Physics	4,800,000
CSC249ADCD05	Mahantesh Halappanavar	Pacific Northwest National Laboratory	1.2.5.3.05 ADCD05-ExaGraph	Computer Science	480,000
CSC249ADCD502	Kenneth Roche	Pacific Northwest National Laboratory	1.2.5.02 ADCD502 Application Assessment	Computer Science	5,800,000
CSC249ADSE08	Steve Hamilton	Oak Ridge National Laboratory	1.2.1.08 ADSE08 ExaSMR	Nuclear Energy	28,000,000
CSC249ADSE03	Andreas Kronfeld	Fermi National Accelerator Laboratory	1.2.1.03 ADSE03-LatticeQCD	Physics	50,400,000
CSC249ADSE09	Paul Kent	Oak Ridge National Laboratory	1.2.1.09 QMCPACK	Materials Science	9,000,000
CSC249ADSE14	Jackie Chen	Sandia National Laboratories	2.2.2.02 ADSE14-Combustion-Pele: Transforming Combustion Science and Technology with Exascale Simulations	Engineering	800,000
CSC249ADSE16	Mark S. Gordon	Iowa State University: Ames Laboratory	2.2.1.03 ADSE16-GAMESS	Chemistry	9,600,000
CSC249ADSE17	Charlie Catlett, Scott Ehling	Argonne National Laboratory	2.2.4.01 Multiscale Coupled Urban Systems	Computer Science	960,000
CSC249ADSE20	Katherine Yelick; Leonid Oliker	Lawrence Berkeley National Laboratory	2.2.4.04 ADSE20-ExaBiome: Exascale Solutions for Microbiome Analysis	Biological Sciences	32,000
CSC250STDM10	Suren Byna	Lawrence Berkeley National Laboratory	1.3.4.10 STDM10-ExaHDF5	Computer Science	24,000
CSC250STMS07	Barry Smith	Argonne National Laboratory	1.3.3.07 STMS07-PETSc/TAO	Mathematics	8,000
CSC250STMS10	Jack Dongarra	University of Tennessee at Knoxville	2.3.3.09 STMS10-SLATE: Software for Linear Algebra Targeting at Exascale	Mathematics	16,000
CSC250STPM11	Jack Dongarra, Earl Carr, George Bosilca, Thomas Herault	University of Tennessee at Knoxville	1.3.1.11 STPM11-Distributed Tasking (ParSEC)	Computer Science	1,800,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CSC250STPM17	Scott Baden	Lawrence Berkeley National Laboratory	1.3.1.17 STPM17-PGAS	Computer Science	16,000
CSC250STTO09	Jack Dongarra, Earl Carr, Heike Jagode, Anthony Danalis	University of Tennessee at Knoxville	1.3.2.09 STTO09-ExaPAPI	Computer Science	3,200,000
CSC250STTO11	John Mellor-Crummey	Rice University	2.3.2.08 STTO11 - Extending HPCToolkit for Exascale	Computer Science	160,000
CWT	Justin M Wozniak	Argonne National Laboratory	Cancer Workflow Toolkit	Biological Sciences	6,000,000
datascience	Venkatram Vishwanath	Argonne National Laboratory	ALCF Data Science and Workflows Allocation	Internal	2,000,000
DD-MacroDFT	Kaushik Bhattacharya	California Institute of Technology	Large-scale *Ab Initio* Simulations of Crystal Defects	Materials Science	2,000,000
DNSTFlow_PostProc	Jonathan Poggie	Purdue University	Direct Numerical Simulation of Compressible, Turbulent Flow	Engineering	5,000,000
DNS_spray	Yue Ling	Baylor University	Direct Numerical Simulation of Interfacial Wave Development and Breakup in a Gas–Liquid Mixing Layer	Engineering	1,280,000
EE-ECP	Valerie Taylor, Xingfu Wu	Argonne National Laboratory	Energy Efficient Trade-off among Execution Time, Power, and Resilience of Two ECP Applications	Computer Science	100,000
DSEM_SUPERSONIC	Farzad Mashayek	University of Illinois at Chicago	Simulation of Supersonic Combustion	Engineering	3,500,000
EarlyPerf_aurora	Scott Parker	Argonne National Laboratory	Porting for Performance: The Software and Programming Environment Project to Enable High Performance Computational Science	Computer Science	5,000,000
el-eff	Eva Zarkadoula	Oak Ridge National Laboratory	2T-MD Model Simulations of High Energy Ion Irradiation	Materials Science	2,200,000
Extreme_Scale_TS	William M. Tang	Princeton Plasma Physics Laboratory	Extreme Scale Turbulence Simulations	Physics	5,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
Enzo-PCello	Michael Norman	University of California, San Diego	Toward Exascale Hydrodynamic Cosmology	Physics	10,000,000
EZ	Franck Cappello	Argonne National Laboratory	Ez: Fast, Effective, Parallel Error-bounded Exascale Lossy Compression for Scientific Data	Computer Science	3,145,728
Full_core_3D_MOC	Kord Smith	Massachusetts Institute of Technology	Full Core PWR Simulation Using 3D Method of Characteristics	Nuclear Energy	11,889,917
HHPMT_PostProc	Karl Hammond, Brian Wirth	The University of Missouri	Plasma-Facing Materials for Fusion Applications	Materials Science	180,000
HiggsTests	Julius Kuti	University of California, San Diego	Code Preparations for the Composite Higgs Model	Physics	1,000,000
HighReyTurb_PostProc	Robert D. Moser	The University of Texas at Austin	Data Analysis of Turbulent Channel Flow at High Reynolds Number	Engineering	6,000,000
HPCTuning	Khaled Ibrahim	Lawrence Berkeley National Laboratory	HPC Applications Tuning Custom Communication Runtime	Computer Science	10,000,000
Hybrid-C-Modelling	Jinxun Liu	U.S. Geological Survey/Western Geographic Science Center	Simulating Global Terrestrial Carbon Sequestration and Carbon Transport to Aquatic Ecosystems – Pilot Study	Earth Science	2,600,000
ICE_BOP	Subramanian Sankaranarayanan	Argonne National Laboratory	Phase Transitions in Water–Ice–Vapor System	Materials Science	74,000,000
ITHALES	Koen Hillewaert	Cenaero	Industrial Turbomachinery High Accuracy LES – iTHALES	Engineering	3,000,000
Injfuel	Marco Arienti	Sandia National Laboratories	High-Fidelity Characterization of Fuel Injection	Engineering	150,000
JCESR	Larry Curtiss	Argonne National Laboratory	Development of High Throughput Methods	Materials Science	10,000,000
keten-composites	Sinan Keten	Northwestern University	Simulation-based Design of Nanocellulose Composites	Materials Science	1,500,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
KineticTurbulence	Jason TenBarge	Princeton University	A Discontinuous Galerkin Fully Kinetic Vlasov Maxwell Study of Plasma Turbulence	Physics	5,000,000
LES-Environment	Trung Bao Le	Medical College of Wisconsin	Large Eddy Simulation of Large-scale Environmental Flows	Biological Sciences	500,000
LMFUSION	Gabriel Font	Lockheed Martin	Simulation of Compact Fusion Reactor in Support of Strategic Energy Assurance and Security	Physics	1,300,000
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	16,000,000
magjet	Edison Liang	Rice University	Magnetized Jet Creation Using a Hollow Ring of Laser Beams	Physics	5,000,000
MagnetismHPC_PostPro	Jonathan Aurnou	University of California, Los Angeles	Frontiers of Planetary and Stellar Magnetism	Earth Science	11,000,000
Maintenance	Mark Fahey	Argonne National Laboratory	LCF Operations System Maintenance	Internal	512,000,000
MDHTProp	Mark Messner	Argonne National Laboratory	Assessing the Scalability of Direct MD Calculations of Structural Material Properties	Engineering	500,000
Motoneuron_NAISE	Charles J. Heckman	Northwestern University Feinberg School of Medicine	Modeling of Motoneuron Activity from Direct Spinal Cord Stimulation	Biological Sciences	4,000,000
MPICH_MCS	Ken Raffenetti	Argonne National Laboratory	MPICH – A High Performance and Widely Portable MPI Implementation	Computer Science	14,000,000
MSAS_DD	Frank X. Vazquez	St. John's University	Multiscale Investigation of a-Synuclein Structural Dynamics	Chemistry	2,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
NCAlloy_TM_Stability	Garritt J. Tucker	Colorado School of Mines	Ascertaining the Thermo-Mechanical Mechanisms of Solute-Stabilized Nanocrystalline Alloys	Materials Science	1,000,000
Nek_VHTR	Masahiro Kawaji	City College of New York	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	32,000,000
PadeOps_inclined-RM	Sanjiva K. Lele	Stanford University	Simulations of Inhomogeneous Shock-Induced Turbulent Mixing	Engineering	15,800,000
PDIFF	Andrey Beresnyak	Naval Research Laboratory	Studying Plasma Diffusion in Line-tied Geometry	Physics	4,000,000
PerfectDefects	Alvaro Vazquez-Mayagoitia	Argonne National Laboratory	Modeling and Simulation of Hafnium Dioxide	Materials Science	2,000,000
Performance	Scott Parker, Ray Loy	Argonne National Laboratory	Performance	Internal	20,000,000
PHASTA_NCSU	Igor A. Bolotnov	North Carolina State University	Multiphase Simulations of Nuclear Reactor Thermal Hydraulics	Engineering	7,000,000
Piston_Bowl_Study	Adam E. Klingbeil	General Electric Company (GE)/Global Research	Numerical Study on the Effect of Spray and Piston Bowl Geometry in Large Bore Diesel Engines	Chemistry	4,000,000
Practical_Strassen	Robert van de Geijn	The University of Texas at Austin	Make Strassens Algorithm Practical	Computer Science	50,000
Quinoa	Jozsef Bakosi	Los Alamos National Laboratory	Asynchronous Navier-Stokes Solver on 3D Unstructured Grids for the Exascale Era	Engineering	2,000,000
radix-io	Philip Carns	Argonne National Laboratory	System Software to Enable Data-intensive Science	Computer Science	1,244,356
ReaxCath	Christopher Knight	Argonne National Laboratory	Reactive Modeling of Battery Cathodes and Interfaces	Chemistry	2,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
rec_sironi	Lorenzo Sironi	Columbia University	Particle-In-Cell Simulations of Explosive Reconnection in Relativistic Magnetically-Dominated Plasmas	Physics	1,500,000
REEs_and_actinides	Deborah A Penchoff	Institute for Nuclear Security	Accelerating Selective Binding of Rare Earth Elements and Actinides	Chemistry	1,936,314
RocketML	Vinay Rao	RocketML	RocketML	Biological Sciences	1,000,000
scalablemoose	Fande Kong	Idaho National Laboratory	MOOSE	Nuclear Energy	1,000,000
SDL_Workshop	Ray Loy	Argonne National Laboratory	ALCF Simulation, Data, and Learning Workshop	Training	5,000,000
SilentOwlFlight_DD	Anupam Sharma	Iowa State University	Unraveling Silent Owl Flight to Develop Ultra-Quiet Energy Conversion Machines	Engineering	5,000,000
slmlearn_mira	Brian Mercer	University of Illinois at Urbana-Champaign/Argonne National Laboratory	Machine Learning for Optimization of Selective Laser Melting Scanning Protocol	Engineering	50,000
SolarWindowsADSP	Jacqueline Cole	University of Cambridge	Data-Driven Molecular Engineering of Solar-powered Windows	Materials Science	112,000,000
SprayWall_UMassD	Roberto Torelli	Argonne National Laboratory/University of Massachusetts-Dartmouth	Direct Numerical Simulations of Droplet-Wall Impingement under High-Pressure Spray Conditions	Engineering	7,200,000
Tools	Scott Parker	Argonne National Laboratory	ALCF Performance Tools	Internal	2,000,000
TopologyMapping	Zhiling Lan	Illinois Institute of Technology	Topology Mapping of Irregular Applications	Physics	250,000
TotalView	Peter Thompson, Ray Loy	Rogue Wave Software, Inc.	TotalView Debugger on Blue Gene P	Internal	100,000
TurbulentLiquidDrop	Arne J. Pearlstein	University of Illinois at Urbana-Champaign	Computation of Transitional and Turbulent Drop Flows for Liquid Carbon Dioxide Drops Rising in Seawater	Engineering	1,000,000

Mira DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
turb_comisso	Luca Comisso	Columbia University	Particle-In-Cell Simulations of Strong Turbulence in Relativistic Plasmas	Physics	1,000,000
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	1,500,000
UrbanExP	Rajeev Jain	Argonne National Laboratory	Urban ECP	Engineering	1,000,000
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	1,000,000
VarRhoFlow	Paul E. Dimotakis	California Institute of Technology	Variable-density Fluid Dynamics	Engineering	2,000,000
Vendor_Support	William E. Allcock	Argonne National Laboratory	Vendor Support	Internal	1,000,000
visualization	Joe Insley, Mike Papka	Argonne National Laboratory	Visualization and Analysis Research and Development for ALCF	Internal	2,000,000
VTR	Dillon Shaver	Argonne National Laboratory	Thermal Hydraulic Simulations for the Versatile Test Reactor	Nuclear Energy	6,000,000
wall_turb_dd	Ramesh Balakrishnan	Argonne National Laboratory	Wall Resolved Simulations of Canonical Wall Bounded Flows	Engineering	18,000,000
XGC_mira	Timothy Williams	Argonne National Laboratory	Testing Modern XGC Code on Mira	Fusion Energy	500,000
				Total Mira DD	**1,368,341,202**

This page intentionally left blank.

Appendix C – Argonne Leadership Computing Facility Director's Discretionary Projects: Theta

January 1, 2018 – December 31, 2018
Director's Discretionary (DD) Projects on 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	Advanced Electronic Structure Methods for Heterogeneous Catalysis and Separation of Heavy Metals	Chemistry	3,000,000
AD_Brain_Imaging	Jiook Cha	Columbia University	Computational Analysis of Brain Connectomes for Alzheimer's Disease	Biological Sciences	10,383,893
ALCF_Getting_Started	Ray Loy	Argonne National Laboratory	ALCF Getting Started	Training	10,000
Allinea	Ray Loy	Argonne National Laboratory	Improved Debugging Memory Usage for Blue Gene/Q (BG/Q)	Internal	100,000
Alzeta	Marc Day	Lawrence Berkeley National Laboratory	Improving Gas Reactor Design with Complex Non-standard Reaction Mechanisms in a Reactive Flow Model	Chemistry	4,700,000
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
ARO_Project	Jonathan Dallmann	Northwestern University	ARO Project Simulation of Turbulent Flow over a Porous Bed	Engineering	200,000
AstroHED_DD	Frederico Fiuza	SLAC National Accelerator Laboratory	Exploring Astrophysical Particle Acceleration in HED Laboratory Plasmas	Physics	10,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
athena_performance	Brian OShea	Michigan State University	Scaling and Performance Enhancement of an Astrophysical Plasma Code	Physics	1,050,000
AtlasADSP	Taylor Childers	Argonne National Laboratory	Advancing the Scalability of LHC Workflows to Enable Discoveries at the Energy Frontier	Physics	45,000,000
atlasMLbjets	Rui Wang	Argonne National Laboratory	Using ML in b-jet Identification at ATLAS	Physics	4,000,000
atlas_aesp	James Proudfoot	Argonne National Laboratory	Simulating and Learning in the ATLAS Detector at the Exascale	Physics	4,000,000
AtomicSolidification	Asle Zaeem	Missouri University of Science and Technology	Large Scale Atomistic Simulations for Predicting Nano/Microstructures and Properties in Solidification of Metals	Materials Science	3,635,440
Atomization	Robert Saye	Lawrence Berkeley National Laboratory	HPC4Mfg: Modeling Paint Behavior During Rotary Bell Atomization	Engineering	7,500,000
ATPESC18_Instructors	Marta Garcia Martinez	Argonne National Laboratory	Argonne Training Program on Extreme Scale Computing for ALL Instructors	Training	100,000
ATPESC2018	Marta Garcia Martinez	Argonne National Laboratory	Argonne Training Program on Extreme-Scale Computing 2018	Computer Science	3,000,000
AuPM6	Emilie Guidez	University of Colorado Denver	Reparametrization of the PM6 Method for Gold Nanoclusters	Chemistry	374,000
AutomatedBench	Kevin Harms	Argonne National Laboratory	Automated Benchmarking	Computer Science	400,000
awa	Nicole Neveu	Argonne National Laboratory	Modeling and Optimization of Advanced Accelerators	Physics	2,000,000
Banhammer_test	Haritha Siddabathuni Som	Argonne National Laboratory	Banhammer Test	Internal	10
Baryons_tesp	Katrin Heitmann	Argonne National Laboratory	Next-Generation Cosmology Simulations with HACC: Challenges from Baryons	Physics	100,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
BioMechGITractSims	Shashank Acharya	Northwestern University	Computational Modeling of the Biomechanics of the Gastrointestinal Tract	Engineering	500,000
BrainImagingADSP	Doga Gursoy	Argonne National Laboratory	Large-scale Computing and Visualization on the Connectomes of the Brain	Biological Sciences	25,000,000
BrainModelingADSP	Felix Schurmann	École Polytechnique Fédérale de Lausanne (EPFL)	Leveraging Non-Volatile Memory Big Data and Distributed Workflow Technology to Leap Forward Brain Modeling	Biological Sciences	10,000,000
Bubblecollapse	Eric Johnsen	University of Michigan	Dynamics and Heat Transfer in Bubble Collapse near Solid Surfaces	Engineering	3,000,000
BuoyantTurbulence	Aleksandr Obabko	Argonne National Laboratory/Finnish Meteorological Institute	Testing a New Theory of Density Effects in Turbulent Transport and Mixing	Earth Science	400,000
candle_aesp	Rick Stevens; Venkat Vishwanath; Thomas Brettin	Argonne National Laboratory	Virtual Drug Response Prediction	Biological Sciences	10,000,000
catalysis_aesp	David Bross	Argonne National Laboratory	Exascale Computational Catalysis	Chemistry	7,000,000
Catalyst	Katherine Riley	Argonne National Laboratory	Catalyst	Internal	10,000,000
CavitationFSI	Kevin Wang	Virginia Polytechnic Institute and State University	High-Fidelity, Fluid-Material Coupled Simulations of Bubble Cloud Cavitation	Engineering	1,500,000
CESM_Highres_Testing	Gerald Meehl, Susan Bates	The National Center for Atmospheric Research	Porting and Benchmarking CESM 1.3 on Theta	Earth Science	3,000,000
cfdml_aesp	Ken Jansen	University of Colorado Boulder	Data Analytics and Machine Learning for Exascale CFD	Engineering	3,000,000
CharmRTS	Laxmikant V. Kale	University of Illinois at Urbana-Champaign	Charm++ and Its Applications	Computer Science	3,999,968
Child_Connectome	Jiook Cha	Columbia University	Developmental Trajectory of Brain and Cognition in Youth in Physiological and Pathological Conditions	Biological Sciences	10,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
climate_severe	Victor Gensini	Northern Illinois University	Anticipating Severe Weather Events via Dynamical Downscaling	Earth Science	3,000,000
Climate_Water	Jiali Wang	Argonne National Laboratory	Linking Climate to Water: Implementing a 4km Regional Climate Model with Hydrologic Model Coupling (WRF-Hydro)	Earth Science	2,304,000
Clinical_Connectome	Jiook Cha	Columbia University	Application of High-throughput Brain Connectomics to Clinical Predictive Modeling	Biological Sciences	5,000,000
CM1-LES	Gokhan Sever	Argonne National Laboratory	Large-Eddy Simulations of Airflow Dynamics and Physics around the Island of Graciosa	Earth Science	1,000,000
CobaltDevel	Paul Rich	Argonne National Laboratory	Cobalt Development	Internal	2,000,000
Coin	Nwike Iloeje, Diane Graziano	Argonne National Laboratory/DMDII	DMDII-Argonne CNC Optimization through Machine Learning	Engineering	100,000
ComEd_VAE	Siby Jose Plathottam	Argonne National Laboratory	Consumer-level Electricity Load Forecasting in Chicagoland Using Unsupervised Deep Learning	Energy Technologies	100,000
CompCatalysis	Rajeev Surendran Assary	Argonne National Laboratory	Accelerated Catalyst Discovery from First Principles Simulations and Machine Learning	Chemistry	5,000,000
CompSpectro_DD	Giulia Galli	The University of Chicago	Electronic and Transport Properties of Complex Interfaces	Materials Science	15,000,000
Comp_Perf_Workshop	Ray Loy	Argonne National Laboratory	Comp_Perf_Workshop	Training	10,000,000
connectomics_aesp	Nicola Ferrier	Argonne National Laboratory	Enabling Connectomics at Exascale to Facilitate Discoveries in Neuroscience	Biological Sciences	10,000,000
CPOX	David Bross	Argonne National Laboratory	Surface Binding Energies for Heterogenous Catalytic Partial Oxidation of Methane	Chemistry	1,155,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CP_SOD1	Shahar Keinan	Cloud Pharmaceuticals, Inc.	Design of Novel Small Molecules to Inhibit SOD1 Aggregation as a Therapy for ALS	Biological Sciences	4,000,000
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	33,000,000
CSC188	Thomas Jackson, Robert Voigt	University of Florida	Demonstration of the Scalability of Programming Environments by Simulating Multi-scale Applications	Engineering	640,000
CSC249ADCD01	Ian Foster	Argonne National Laboratory	1.2.5.3.01 ADCD01-CODAR	Computer Science	201,600
CSC249ADCD02	Timothy Germann	Los Alamos National Laboratory	2.2.6.04 ADCD02-COPA: Co-Design Center for Particle Applications	Physics	1,280,000
CSC249ADCD04	Tzanio Kolev, Misun Min, Paul Fischer	Lawrence Livermore National Laboratory	1.2.5.3.04 Center for Efficient Exascale Discretizations (CEED)	Computer Science	281,088
CSC249ADCD05	Mahantesh Halappanavar	Pacific Northwest National Laboratory	1.2.5.3.05 ADCD05-ExaGraph	Computer Science	1,280,000
CSC249ADCD502	Kenneth Roche	Pacific Northwest National Laboratory	1.2.5.02 ADCD502 Application Assessment	Computer Science	2,936,000
CSC249ADCD504	David Richards	Lawrence Livermore National Laboratory	2.2.6.01 ADCD504-Proxy Applications	Computer Science	12,800
CSC249ADOA01	Rick L. Stevens	Argonne National Laboratory	1.2.3.01 ADOA01 CANDLE	Biological Sciences	36,941,248
CSC249ADSE01	Salman Habib	Argonne National Laboratory	1.2.1.01 ADSE01-ExaSky	Physics	6,760,000
CSC249ADSE03	Andreas Kronfeld	Fermi National Accelerator Laboratory	1.2.1.03 ADSE03-LatticeQCD	Physics	4,480,000
CSC249ADSE04	Arthur F. Voter	Los Alamos National Laboratory	1.2.1.04 ADSE04-EXAALT – Molecular Dynamics at the Exascale	Nuclear Energy	2,304,000
CSC249ADSE08	Steve Hamilton	Oak Ridge National Laboratory	1.2.1.08 ADSE08 ExaSMR	Nuclear Energy	3,840,000
CSC249ADSE09	Paul Kent	Oak Ridge National Laboratory	1.2.1.09 QMCPACK	Materials Science	1,640,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CSC249ADSE14	Jackie Chen	Sandia National Laboratories	2.2.2.02 ADSE14-Combustion-Pele: Transforming Combustion Science and Technology with Exascale Simulations	Engineering	12,800,000
CSC249ADSE16	Mark S. Gordon	Iowa State University	2.2.1.03 ADSE16-GAMESS	Chemistry	1,344,000
CSC249ADSE17	Charlie Catlett, Scott Ehling	Argonne National Laboratory	2.2.4.01 Multiscale Coupled Urban Systems	Computer Science	96,000
CSC249ADSE20	Katherine Yelick; Leonid Oliker	Lawrence Berkeley National Laboratory	2.2.4.04 ADSE20-ExaBiome: Exascale Solutions for Microbiome Analysis	Biological Sciences	960,000
CSC249ADSE22	Zhenyu Huang	Pacific Northwest National Laboratory	1.2.1.22 ADSE22-ExaSGD	Energy Technologies	3,304,000
CSC250STDM10	Suren Byna	Lawrence Berkeley National Laboratory	1.3.4.10 STDM10-ExaHDF5	Computer Science	192,000
CSC250STDM11	Scott Klasky	Oak Ridge National Laboratory	2.3.4.09 STDM11-ADIOS Framework for Scientific Data on Exascale Systems	Computer Science	2,880,000
CSC250STDM12	Rob Ross	Argonne National Laboratory	2.3.4.10 STDM12-DataLib: Data Libraries and Services Enabling Exascale Science	Computer Science	64,000
CSC250STMS05	Lois Curfman McInnes	Argonne National Laboratory	2.3.3.01 STMS05-Extreme-scale Scientific xSDK for ECP	Mathematics	1,280,000
CSC250STMS07	Barry Smith	Argonne National Laboratory	1.3.3.07 STMS07-PETSc/TAO	Mathematics	1,280,000
CSC250STMS10	Jack Dongarra	University of Tennessee at Knoxville	2.3.3.09 STMS10-SLATE: Software for Linear Algebra Targeting at Exascale	Mathematics	384,000
CSC250STMS11	Jack Dongarra	University of Tennessee at Knoxville	2.3.3.10 STMS11-PEEKS: Production-ready, Exascale-enabled Krylov Solvers for Exascale Computing	Mathematics	640,000
CSC250STPM07	Bruce Palmer	Pacific Northwest National Laboratory	1.3.1.07 STPM07 xGA	Computer Science	320,000
CSC250STPM11	Jack Dongarra, Earl Carr, George Bosilca, Thomas Herault	University of Tennessee at Knoxville	1.3.1.11 STPM11-Distributed Tasking (ParSEC)	Computer Science	640,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CSC250STPM17	Scott Baden	Lawrence Berkeley National Laboratory	1.3.1.17 STPM17-PGAS	Computer Science	102,400
CSC250STTO09	Jack Dongarra, Earl Carr, Heike Jagode, Anthony Danalis	University of Tennessee at Knoxville	1.3.2.09 STTO09-ExaPAPI	Computer Science	5,480,000
CSC250STTO10	Mary Hall	University of Utah	2.3.2.07 STTO10-Autotuning Compiler Technology for Cross-Architecture Transformation and Code Generation	Computer Science	448,000
CSC250STTO11	John Mellor-Crummey	Rice University	2.3.2.08 STTO11 - Extending HPCToolkit for Exascale	Computer Science	640,000
CSSI	Wei Jiang	Argonne National Laboratory	Scalable Reconstruction of X-ray Scattering Imaging for Nanomaterials	Physics	6,000,000
CWT	Justin M Wozniak	Argonne National Laboratory	Cancer Workflow Toolkit	Biological Sciences	1,000,000
darkskyml_aesp	Salman Habib	Argonne National Laboratory	Dark Sky Mining	Physics	2,000,000
datascience	Venkatram Vishwanath	Argonne National Laboratory	ALCF Data Science and Workflows Allocation	Internal	110,000,000
DataSpaces	Manish Parashar	Rutgers University	DataSpaces Data Management Framework	Computer Science	100,000
DD-ALSDFT	Kaushik Bhattacharya	California Institute of Technology	Accelerated Linear Scaling Density Functional Theory: Simulating Defects in Crystalline Magnesium	Materials Science	5,000,000
DDTM-NE	Rui Hu	Argonne National Laboratory	Data-Driven Turbulence Modeling for Nuclear Energy Applications	Nuclear Energy	100,000
DE-SC0017033	Yueqin Huang	Cyentech Consulting LLC	Hardening of DOE UQ Codes for Ultra-Deep Proactive Geosteering in Oilfield Service	Earth Science	120,000
DFT-FE	Phani Motamarri, Vikram Gavini	University of Michigan	Large-scale Real-space Electronic Structure Calculations for Understanding Energetics of Complex Defects in Materials	Materials Science	4,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
DistrElStructCalcs	Marco Govoni	Argonne National Laboratory	For Everyone A21: Distributed Electronic Structure Calculations Using a Globus-Enabled Programmable Infrastructure	Materials Science	15,000,000
Distributed_Learning	Eliu Huerta	University of Illinois at Urbana-Champaign	Deep Learning at Scale for Gravitational Wave Astrophysics	Physics	3,000,000
DLHMC	James Osborn	Argonne National Laboratory	Deep Learning HMC	Physics	1,000,000
DMH	Xian-He Sun	Illinois Institute of Technology	Utilizing Memory Parallelism for High Performance Data Processing	Computer Science	20,000
DMultiFMSI_DD	Aaron Packman	Northwestern University/University of Chicago/ Lawrence Livermore National Laboratory	Multiscale Investigation of Fluid-Materials-Solute Interactions for Water Treatment	Engineering	2,000,000
DNS_wavyWall	Sparsh Ganju	University of Kentucky	Direct Numerical Simulations of Turbulent Channel Flows with Sinusoidally Distributed Roughness	Engineering	1,000,000
DOECyberDef_Comp	Jeff Neel	Argonne National Laboratory	DOE Cyber Def Competition	Internal	2,000
DowAcclimation	Wililam J. Edsall	The Dow Chemical Company	Dow HPC Team Acclimating to Theta	Chemistry	50,000
duanl	Lian Duan	Missouri University of Science and Technology	Numerical Simulation of Acoustic Radiation from High-Speed Turbulent Boundary Layers	Engineering	1,000,000
DUNE-LBNF	Nikolai Mokhov	Fermi National Accelerator Laboratory	MARS Energy Deposition and Neutrino Flux Simulations	Physics	20,000,000
EarlyPerf_aurora	Scott Parker	Argonne National Laboratory	Porting for Performance: The Software and Programming Environment Project to Enable High Performance Computational Science	Computer Science	6,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
ecp-testbed-01	Rajeev Thakur, Julia White, Doug Kothe, Pat McCormick, Susan Coghlan	Argonne National Laboratory	Exascale Computing Project (ECP) Testbed	Computer Science	41,415,598
ECP_SDK	Sameer Shende	University of Oregon	Deploying the ECP SDK Software Stack at ALCF	Computer Science	100,000
EE-ECP	Valerie Taylor, Xingfu Wu	Argonne National Laboratory	Energy Efficient Trade-off among Execution Time, Power, and Resilience of Two ECP Applications	Computer Science	1,100,000
el-eff	Eva Zarkadoula	Oak Ridge National Laboratory	2T-MD Model Simulations of High Energy Ion Irradiation	Materials Science	550,000
Enzyme_Sim	Tao Li	Argonne National Laboratory	Stability and Activity Simulations of Biomass Enzyme under Complex Solvent Environment	Biological Sciences	10,000,000
eso-sim	Wenjun Kou	Northwestern University	Biophysical Modeling and Data Analysis on Esophageal Physiology	Biological Sciences	1,000,000
EWN-DP-Tool	Getnet Betrie	Argonne National Laboratory	A Tool for Energy-Water-Nexus Analysis: A Deep Learning Framework	Earth Science	1,400,000
ExaChemPrep	Colleen Bertoni	Argonne National Laboratory	Benchmarking and Characterization of Electronic Structure Codes in Preparation for Exascale	Chemistry	700,000
EXASTEEL-THETA	Axel Klawonn, Oliver Rheinbach	TU Bergakademie Freiberg	Simulation of Dual-Phase Steel on Theta	Engineering	1,500,000
EXATHERM	David Bross	Argonne National Laboratory	Automated Generation of Customized Quantum Chemical Composite Methods Validated via the ATcT Database	Chemistry	1,000,000
ExaWindFarm_DD	Anouar Benali	Argonne National Laboratory	Wind Farm Exascale	Energy Technologies	500,000
ExM	Justin M Wozniak	Argonne National Laboratory	Extreme Many-task Computing with Swift	Computer Science	25,000
FFTBench	William Scullin	Argonne National Laboratory	FFT Benchmarking for Imaging Applications	Computer Science	3,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
FLoRIN	Walter Scheirer	University of Notre Dame	Large-Scale Segmentation and Reconstruction of Neural Microscopy	Biological Sciences	3,000,000
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
Full_Core_TRRM	John Tramm	Massachusetts Institute of Technology	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	First-Principles Simulations of Functional Materials for Energy Conversion	Materials Science	90,000,000
fusiondl_aesp	William Tang	Princeton University	Accelerated Deep Learning Discovery in Fusion Energy Science	Fusion Energy	1,000,000
fvGFS	Jeffrey Durachta	Geophysical Fluid Dynamics Laboratory (GFDL)	Global Forecast System Using the Finite-Volume Cubed-Sphere Dynamical Core and GFS Physics	Earth Science	500,000
G3DNet	Taylor Childers	Argonne National Laboratory	3D GraphCNN for Particle Detectors	Computer Science	2,000,000
GA-bfg3d	Mark Kostuk	General Atomics	General Atomics Next Generation Skew-Symmetric Fluid Solver	Physics	800,000
GasJetsCylPrio_tesp	Christos Frouzakis	ETH Zurich	Flow, Mixing, and Combustion of Transient Turbulent Gaseous Jets in Confined Cylindrical Geometries	Chemistry	10,000,000
GAtor_scaling	Noa Marom	Carnegie Mellon University	Scaling Benchmarks for the GAtor CSP Code	Materials Science	1,000,000
GFDL_Ensemble	Thomas Robinson	Geophysical Fluid Dynamics Laboratory (GFDL)	Ensemble-based Regression Tests for Climate Model Reproducibility	Earth Science	1,100,000
HACC_aesp	Katrin Heitmann	Argonne National Laboratory	Extreme-Scale Cosmological Hydrodynamics	Physics	2,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
HadronicLightTheta	Christoph Lehner	Brookhaven National Laboratory	Hadronic Light-by-light Scattering and Vacuum Polarization Contributions to the Muon g-2 from LQCD - Finite-volume Studies	Physics	20,000,000
HEP_on_HPC	Jim Kowalkowski	Fermi National Accelerator Laboratory	HEP Analysis Workflows on HPC	Physics	1,200,000
hifiturbfsi	Ivan Bermejo-Moreno	University of Southern California (USC)	High-fidelity Simulation of Supersonic Turbulent Flow-structure Interaction and Mixing	Engineering	6,000,000
HiResComb	Gabriel Staffelbach	Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique (CERFACS)	High Resolution LES for Combustion	Chemistry	900,000
HPC4Mfg_ACS	Deborah Bard	Lawrence Berkeley National Laboratory	Accelerating the Industrial Application of Energy-Efficient Chemical Separation	Materials Science	9,500,000
HyperparADSP	Pierre Baldi	University of California, Irvine	Massive Hyperparameter Searches on Deep Neural Networks Using Leadership Systems	Computer Science	10,000,000
IBR2M	Alberto Talamo	Argonne National Laboratory	Transient Simulation of IBR2M	Nuclear Energy	3,000,000
ICE_Nek5000	Muhsin M Ameen	Argonne National Laboratory	Scalable Internal Combustion Engine Simulations Using Nek5000	Energy Technologies	5,000,000
IL_elec	Zhengcheng Zhang	Argonne National Laboratory	Ionic Liquid as a Potential Electrolyte of High Performance Lithium Ion Battery	Chemistry	1,000,000
InSituVis2018	Joseph Insley	Argonne National Laboratory	In Situ Visualization Hackathon 2018	Computer Science	500,000
Intel	Kalyan Kumaran	Argonne National Laboratory	Intel Employees in Support of Theta	Internal	10,000,000
IntelPyTorch	Elise Jennings	Argonne National Laboratory	Intel Optimized PyTorch for XC40 Systems	Physics	100,000
interconnect_bench	Sudheer Chunduri	Argonne National Laboratory	Collaboration with Cray on Interconnect-related Studies	Computer Science	200,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
JCESR	Larry Curtiss	Argonne National Laboratory	Development of High Throughput Methods	Materials Science	5,000,000
Job_Interference	Zhiling Lan	Illinois Institute of Technology	Workload Interference Analysis on Theta	Computer Science	1,500,000
K-Ras-Model	Sunhwan Jo	Qulab, Inc	Dynamic Model of K-Ras WT and Mutant	Biological Sciences	1,000,000
KineticTurbulence	Jason TenBarge	Princeton University	A Discontinuous Galerkin Fully Kinetic Vlasov Maxwell Study of Plasma Turbulence	Physics	2,000,000
Landscapes	Michael Prentiss	University of Cambridge	Exploring Protein Folding Energy Landscapes	Biological Sciences	1,000,000
large3dxrayADSP	Chris Jacobsen	Northwestern University	X-ray Microscopy of Extended 3D Objects: Scaling toward the Future	Biological Sciences	8,000,000
LatticeQCD_aesp	Paul Mackenzie	Fermi National Accelerator Laboratory	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
load_modelling	Siby Jose Plathottam	Argonne National Laboratory	Consumer-level Electricity Load Forecasting Using Unsupervised Deep Learning	Energy Technologies	100,000
LPI_shlgnt	Jun Li	University of California, San Diego	Hot Electron Scaling and Energy Coupling in Nonlinear Laser Plasma Interactions	Physics	2,000,000
lqcdml_aesp	William Detmold	Massachusetts Institute of Technology	Machine Learning for Lattice Quantum Chromodynamics	Physics	6,000,000
LQCD_VeloC	Chulwoo Jung	Brookhaven National Laboratory	Lattice QCD with VeloC	Physics	1,200,000
LSAN0001	Juan R. Perilla	University of Delaware	Mechanochemical Mechanism of Vesicle Scission by MxB and Dynamin	Biological Sciences	1,000,000
LSSTADSP_DESC	Katrin Heitmann	Argonne National Laboratory	Realistic Simulations of the LSST Survey at Scale – DESC	Physics	53,000,000
LSSTADSP_HACC	Katrin Heitmann	Argonne National Laboratory	Realistic Simulations of the LSST Survey at Scale – HACC	Physics	27,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
magjet	Edison Liang	Rice University	Magnetized Jet Creation Using a Hollow Ring of Laser Beams	Physics	2,000,000
magstructsADSP	Efthimios Kaxiras, Trevor David Rhone	Harvard University	Machine Learning Magnetic Properties of van der Waals Heterostructures	Materials Science	10,000,000
Maintenance	Mark Fahey	Argonne National Laboratory	LCF Operations System Maintenance	Internal	10,000,000
MassStarII	Yan-Fei Jiang	University of California, Santa Barbara	Radiation Magneto-Hydrodynamic Simulations of Massive Stars	Physics	5,000,000
matml_aesp	Noa Marom	Carnegie Mellon University	Many-Body Perturbation Theory Meets Machine Learning to Discover Singlet Fission Materials	Materials Science	9,000,000
MCPSI_DD	Ahren Jasper	Argonne National Laboratory	High Temperature Kinetics and Dynamics of Fluxional Molecules	Chemistry	2,500,000
MFIX_VVUQ	Aytekin Gel	ALPEMI Consulting	Development of a UQ Workflow for CFD-DEM Simulations Pulsatile Fluidized Beds with Experimental Data	Engineering	125,000
MIO_HP_DD	Noa Marom	Carnegie Mellon University	Materials and Interfaces for Organic and Hybrid Photovoltaics	Materials Science	1,000,000
ML4LGT	Samuel Foreman	Argonne National Laboratory	ML for Lattice Gauge Theory	Physics	1,000,000
MLResilience	Yanjing Li	The University of Chicago	Machine Learning Guided Cross-Layer Resilience	Computer Science	1,000,000
mmaADSP	Eliu Huerta	University of Illinois at Urbana-Champaign	Deep Learning at Scale for Multimessenger Astrophysics through the NCSA-Argonne Collaboration	Physics	10,000,000
MOAB_App	Vijay Mahadevan	Argonne National Laboratory	MOAB Algorithmic Performance Portability	Mathematics	500,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
Monochalcogenides	Salvador Barraza-Lopez	University of Arkansas, Fayetteville	Structural Phase Transitions and Catalytic Properties of Group-IV Monochalcogenide Monolayers	Materials Science	1,000,000
Motoneuron_NAISE	Charles J. Heckman	Northwestern University	Modeling of Motoneuron Activity from Direct Spinal Cord Stimulation	Biological Sciences	1,500,000
MPICH_MCS	Ken Raffenetti	Argonne National Laboratory	MPICH – A High Performance and Widely Portable MPI Implementation	Computer Science	16,000,000
MSAS_DD	Frank X. Vazquez	St. John's University	Multiscale Investigation of a-Synuclein Structural Dynamics	Chemistry	2,000,000
MueLu	Christopher Siefert	Sandia National Laboratories, California	Algebraic Multigrid in Trilinos	Mathematics	1,000,000
multiphysics_aesp	Amanda Randles	Duke University/Oak Ridge National Laboratory	Extreme-scale In Situ Visualization and Analysis of Fluid-Structure-Interaction Simulations	Engineering	1,000,000
Multirefer-QMC	Anouar Benali	Argonne National Laboratory	Multireference Trial Wavefunction for Quantum Monte Carlo	Chemistry	2,000,000
MultiscaleADSP	Rathakrishnan Bhaskaran, Marc Edgar	General Electric Company (GE)/Global Research	Enabling Multi-Scale Physics for Industrial Design using Deep Learning Networks	Engineering	16,000,000
MultSimHem_dd	George Karniadakis	Brown University	Multiscale Simulations of Hematological Disorders	Biological Sciences	2,000,000
n2ase	Graham Fletcher	Argonne National Laboratory	Valence Bond Study of Nitrogenase	Chemistry	2,000,000
NAMD_aesp	Benoit Roux	The University of Chicago	Free Energy Landscapes of Membrane Transport Proteins	Biological Sciences	2,000,000
NAQMC_RMD_aesp	Aiichiro Nakano	University of Southern California	Metascalable Layered Materials Genome	Materials Science	24,000,000
NearSunTurb	Jean C Perez	Florida Institute of Technology	Plasma Turbulence near the Sun	Physics	2,000,000
NekBison	Ronald Rahaman	Argonne National Laboratory	Nek-Bison Coupling for NEAMS Toolkit Integration	Nuclear Energy	500,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
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	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	100,000
NetworkScience	Mehmet H Gunes	University of Nevada, Reno	Analysis of Complex Networks on HPC: A Case Study of the Internet	Computer Science	900,000
NWChemEx_aesp	Thomas Dunning, Alvaro Mayagoitia	Princeton Plasma Physics Laboratory	NWChemEx: Tackling Chemical, Materials, and Biochemical Challenges in the Exascale Era	Chemistry	10,000,000
OF_ICING	Roberto Paoli	University of Illinois at Chicago	Aircraft Icing Simulations with OpenFoam	Engineering	2,000,000
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	5,000
PCI_MPS	Miaoqi Chu	Argonne National Laboratory	Phase Contrast Imaging Multiphysics Simulation	Physics	1,000,000
PCI_WP	Miaoqi Chu	Argonne National Laboratory	Imaging Simulations for APS Fuel Spray Experiments	Physics	1,000,000
Performance	Scott Parker, Ray Loy	Argonne National Laboratory	Performance	Internal	20,000,000
perf_research	Sudheer Chunduri	Argonne National Laboratory	Performance group external facing research	Computer Science	200,000
PerNitro	David Cooper	University of Liverpool	Performance of Nitrogenase VSVB Calculations	Chemistry	2,000,000
PETScKrylov	Hannah Morgan, Todd Munson	Argonne National Laboratory	The Influence of Noise on Krylov Methods	Computer Science	250,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
PHASTA_aesp	Kenneth Jansen	University of Colorado Boulder	Extreme Scale Unstructured Adaptive CFD: From Multiphase Flow to Aerodynamic Flow Control	Engineering	7,000,000
PHASTA_AF	Jun Fang	Argonne National Laboratory	Detached Eddy Simulations of Separated Flows	Engineering	5,000,000
pieiclat	Ian Cloet	Argonne National Laboratory	Pion and Kaon Quark-Gluon Structure from Lattice QCD	Physics	2,000,000
Pipe_Transition_DNS	Xiaohua Wu	Royal Military College of Canada	Discovering the Transition Mechanism in 8000 Radii Long Pipe Flow without Axially Periodic Boundary Condition	Engineering	2,000,000
PKrylovPetsc	Oana Marin	Argonne National Laboratory	PETSc Pipeline Krylov	Computer Science	1,000,000
Polymersomes	Abelardo Ramirez-Hernandez	The University of Chicago	From Complex Macromolecules to Kinetically Arrested Hierarchical Assemblies	Chemistry	1,000,000
PrincetonConnectome	Sebastian Seung	Princeton University	Petascale Neural Circuit Reconstruction	Biological Sciences	25,000
PROTEUS_KNL	Yeon Sang Jung	Argonne National Laboratory	Porting High Fidelity Neutron Transport Code PROTEUS-MOC to KNL Platform	Nuclear Energy	1,000,000
PSFMat	Paul Kent	Oak Ridge National Laboratory	Predictive Simulations of Functional Materials	Materials Science	32,000,000
qcarchive	Daniel G. A. Smith	The Molecular Sciences Software Institute	The MolSSI Quantum Chemistry Archive Project	Chemistry	1,000,000
QCSim	Hal Finkel	Argonne National Laboratory	Simulating Realistic Quantum Computers	Computer Science	2,400,000
QMC-EFMO	Federico Zahariev	Iowa State University/Ames Laboratory	Linearly Scaling QMC through EFMO Fragmentation	Chemistry	3,000,000
QMC-SSD	Ye Luo	Argonne National Laboratory	Evaluate the Use of Local SSDs as Part of the Memory Hierarchy for QMCPACK	Materials Science	500,000
QMCCorrElec	Huihuo Zheng	Argonne National Laboratory	Quantum Monte Carlo Modeling of Strongly Correlated Electronic Systems	Physics	2,500,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
QMCPACK_aesp	Anouar Benali	Argonne National Laboratory	Extending Moore's Law Computing with Quantum Monte Carlo	Materials Science	65,000,000
QMC_ankit	Ankit Mahajan	University of Colorado Boulder	Study of Strongly Correlated Electrons	Chemistry	500,000
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	1,100,000
Quinoa	Jozsef Bakosi	Los Alamos National Laboratory	Asynchronous Navier-Stokes Solver on 3D Unstructured Grids for the Exascale Era	Engineering	2,000,000
radix-io	Philip Carns	Argonne National Laboratory	System Software to Enable Data-intensive Science	Computer Science	5,000,000
ResourceAllocation	Michael T Tamillow	Unicorn Markets	Artificial Intelligence Tools for Allocation Problems	Biological Sciences	1,100,000
RL-fold	Arvind Ramanathan	Oak Ridge National Laboratory	Targeting Intrinsically Disordered Proteins Using Artificial Intelligence-driven Molecular Simulations	Biological Sciences	1,000,000
SBC-UBEM	Ralph Muehleisen	Argonne National Laboratory	Scalable Bayesian Calibration of Urban-Scale Building Energy Models	Energy Technologies	200,000
SC18SCC	William Scullin	Illinois Institute of Technology	IIT Student Cluster Challenge Preparation: 2018	Computer Science	500,000
scalablemoose	Fande Kong	Idaho National Laboratory	MOOSE	Nuclear Energy	2,000,000
SCO_thermal_conduct	Bilge Yildiz	Massachusetts Institute of Technology	Thermal Conductivity Modulated by Oxygen and Hydrogen Ions in SrCoOx	Materials Science	200,000
SDL_Workshop	Ray Loy	Argonne National Laboratory	ALCF Simulation, Data, and Learning Workshop	Training	20,100,000
SENSEI	Venkat Vishwanath	Argonne National Laboratory	Scalable Analysis Methods and *In Situ* Infrastructure for Extreme Scale Knowledge Discovery	Computer Science	3,000,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
Slip-Flows	Ratnesh Shukla	Indian Institute of Science	DNS of Flow Past Hydrodynamically Slipping Surfaces	Engineering	1,500,000
slmlearn_mira	Brian Mercer	University of Illinois at Urbana-Champaign/Argonne National Laboratory	Machine Learning for Optimization of Selective Laser Melting Scanning Protocol	Engineering	50,000
SolarScaling	Nicholas Featherstone	University of Colorado Boulder	Performance Analysis of Rayleigh on Theta	Physics	3,000,000
SolarWindowsADSP	Jacqueline Cole	University of Cambridge	Data-Driven Molecular Engineering of Solar-powered Windows	Materials Science	55,000,000
SPS	Vadim Roytershteyn	Space Science Institute	Spectral Kinetic Simulations of Plasmas	Physics	640,000
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
TomoDev	William Scullin	Argonne National Laboratory	Tomographic Software Development	Computer Science	750,000
Tools	Scott Parker	Argonne National Laboratory	ALCF Performance Tools	Internal	5,000,000
TotalView	Peter Thompson, Ray Loy	Rogue Wave Software, Inc.	TotalView Debugger on Blue Gene P	Internal	100,000
Transonic_Buffet	Kai Kruger Bastos	Duke University	The Buffet Phenomenon in Subsonic, Transonic and Hypersonic Flow Regimes	Engineering	2,406,400
TriboChem_DD	Ashlie Martini	University of California, Merced	Reactive Molecular Dynamics Simulations of Protective Film Formation in Sliding Interfaces	Materials Science	1,000,000
TurbLowBeta	Luca Franci	Queen Mary University of London	Hybrid Simulations of Kinetic Plasma Turbulence at Low Electron Beta	Physics	7,256,478
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	4,049,136

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
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	9,000,000
ultrafastxrayADSP	Jin Wang	Argonne National Laboratory	Developing High-Fidelity Dynamic and Ultrafast X-Ray Imaging Tools for APS-Upgrade	Physics	8,000,000
UMass_MultiPhaseFlow	David Schmidt	University of Massachusetts-Amherst	Fuel Spray Modeling and Simulation of Wave Loads on Offshore Wind Turbines	Engineering	3,226,707
UrbanExP	Rajeev Jain	Argonne National Laboratory	Urban ECP	Engineering	1,500,000
UrbanLES	Robert Jacob	Argonne National Laboratory	LES Simulations of the Urban Boundary Layer	Earth Science	1,000,000
VariantToMinibrain	Win van Drongelen	The University of Chicago	Computational Selection of Genetic Variants Underlying Epilepsy	Biological Sciences	200,000
Vehicle_Charging_Sim	Kaizhong Gao	Argonne National Laboratory	Vehicle Charging Model for Extreme Fast Charging Process	Energy Technologies	500,000
VeloC	Bogdan Nicolae	Argonne National Laboratory	VeloC: Very Low Overhead Checkpointing System	Computer Science	10,000
visualization	Joe Insley, Michael E. Papka	Argonne National Laboratory	Visualization and Analysis Research and Development for ALCF	Internal	2,000,000
Viz_Support	Joe Insley	Argonne National Laboratory	Visualization Support	Computer Science	2,000,000
WASH123D	Cheng Wang	Argonne National Laboratory	Physics-based Modeling of Complex Water System	Earth Science	250,000
WaterHammer	Hong Zhang	Argonne National Laboratory	Water Hammer Simulation	Mathematics	20,000
WGSanalysis	Elizabeth McNally	Northwestern University	Large-scale Alignment and Analysis of Whole Human and Mouse Genomes, with Focus on Realigning to HG 38 and Harmonized Workflows	Biological Sciences	200,000

Theta DD	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
XFELO	Henry Freund	University of New Mexico/Calabazas Creek Research	Simulation of x-ray FEL Oscillators	Physics	2,050,000
XGC_aesp	C.S. Chang	Princeton Plasma Physics Laboratory	High-fidelity Simulation of Fusion Reactor Boundary Plasmas	Fusion Energy	3,000,000
yakushin	Igor Yakushin	The University of Chicago	Using Deep Learning for Gravitational Lens Classification and Parameter Regression	Physics	100,000
				Total Theta DD	**1,524,509,766**

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

Disclaimer
This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor UChicago Argonne, LLC, nor any of their employees or officers, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of document authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, Argonne National Laboratory, or UChicago Argonne, LLC.

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.