Argonne NATIONAL LABORATORY | Argonne Leadership Computing Facility

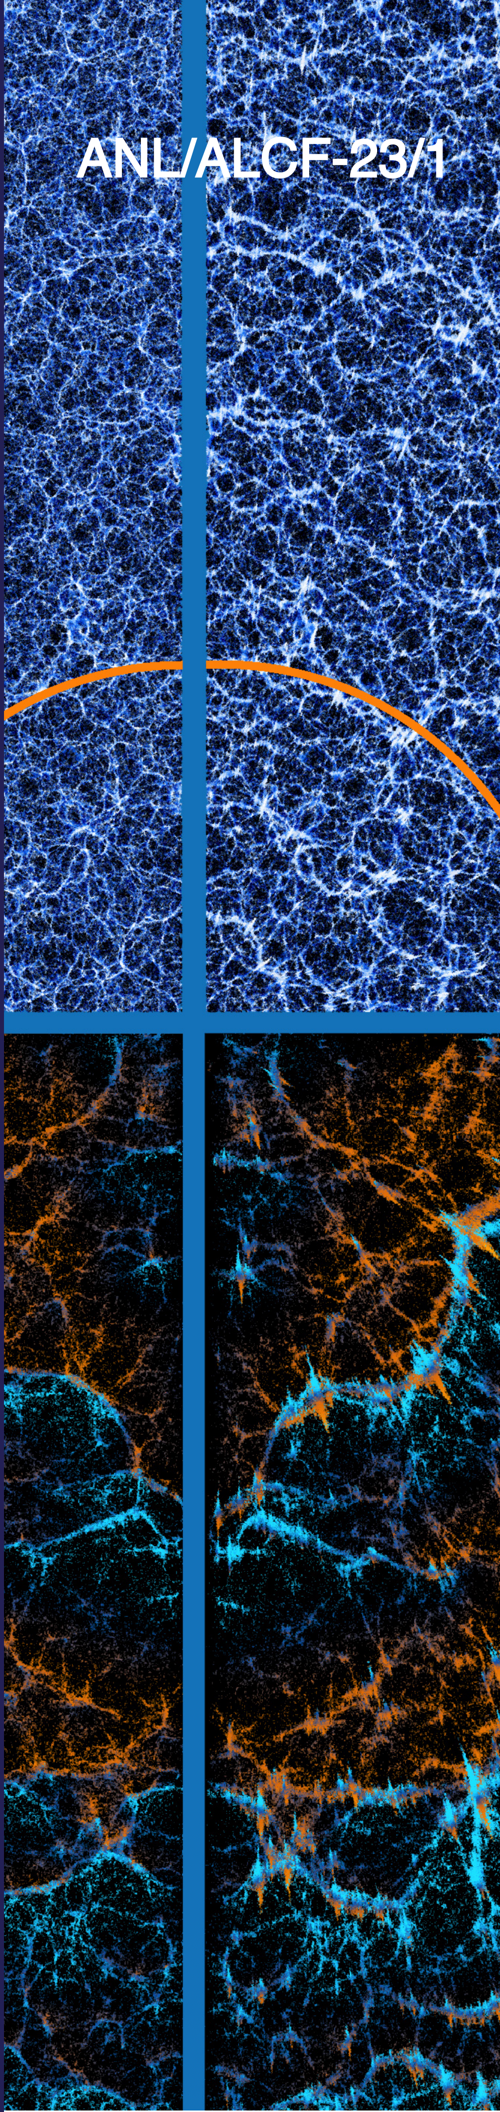

ANL/ALCF-23/1

2022 Operational Assessment Report

On the cover: *Visualization of particle light cone data from Farpoint, a high-resolution cosmology simulation at the gigaparsec scale. The color in the upper panels represents depth in the disk of particles, whereas the colors in the lower panels represent line-of-sight velocities. The left panels render the co-moving particle positions while the right panels demonstrate the so-called Fingers of God effect from redshift space distortion.*

Image credit: *ALCF's Visualization and Data Analysis team and the HACC team.*

Contents

Background xi

Executive Summary ES-1

Response to Recommendation from the PY OAR RTR-1

Section 1. User Support Results 1-1

ALCF Response 1-1
    Survey Approach 1-2
        Likert Scale and Numeric Mapping 1-2
1.1 User Support Metrics 1-3
1.2 Problem Resolution Metrics 1-4
1.3 User Support and User Engagement 1-5
    1.3.1 Tier 1 Support 1-5
        1.3.1.1 Phone, Slack, and E-mail Support 1-5
        1.3.1.2 User Account Management Software (Userbase3) Improvement 1-6
            Significant Improvements in Automated Testing for Userbase3 1-6
            Improvements in Process for Access Control and NDA Projects in Userbase3 1-7
            Support for AI Testbeds and Polaris in Userbase3 1-7
            Auto Subscribing Users to System Mailing Lists 1-7
        1.3.1.3 Performance Portability Study 1-7
        1.3.1.4 Gaussian Process Modeling at HPC Scales 1-7
    1.3.2 Application Support for Individual Projects 1-8
            Enabling analysis of WRF outputs for ALCC project 1-8
            Initial porting of LAMMPS code to GPUs as part of DD project for future ALCC Submission 1-8
    1.3.3 User Engagement 1-8
        1.3.3.1 User Advisory Council 1-8
        1.3.3.2 Training Activities 1-8
            ALCF GPU Hackathon (virtual) 1-8
            Computational Performance Workshop (virtual) 1-9
            2022 INCITE Proposal Writing Webinars (virtual) 1-9
            2022 NVIDIA ALCF GPU Hackathon (virtual) 1-9
            Argonne Training Program on Extreme-Scale Computing (ATPESC) 1-9
            Simulation, Data, and Learning (SDL) Workshop (virtual) 1-9
            Intro to AI-driven Science on Supercomputers: A Student Training Series (virtual) 1-9
            Aurora Early Science Program (ESP) Dungeons, Hackathons, and Workshops (virtual) 1-10
            ALCF Webinars (virtual) 1-10
            AI/Machine Learning (ML) Workshops/Training 1-11
        1.3.3.3 Community Outreach 1-11
            STEM Activities 1-11

Contents (Cont.)

Summer 2022 Research Internships ... 1-11
2022 Introduce a Girl to Engineering Day ... 1-12
2022 Science Careers in Search of Women ... 1-12
2022 CodeGirls @ Argonne Camp (virtual) ... 1-12
Argonne-NIU AI Camp ... 1-12
2022 Coding for Science Camp ... 1-12
2022 Big Data Camp ... 1-12
Hour of Code ... 1-13
1.3.3.4 General Outreach ... 1-13
Diversity, Equity, and Inclusion ... 1-13
Facility Tours ... 1-13
1.3.4 Communications ... 1-13
1.3.4.1 Website Support Center Continuous Improvement ... 1-13
1.3.4.2 Expanding User Experience Across the ALCF ... 1-14
1.3.4.3 Consistent Cadence of ALCF Impact on Exascale and AI Efforts ... 1-14
1.3.4.4 Communicating Scientific Impact ... 1-14
1.3.4.5 Messaging for Users and Community ... 1-15
Conclusion ... 1-16

Section 2. Operational Performance ... 2-1

2.1 ALCF Response ... 2-1
2.2 ALCF Production Resources Overview ... 2-3
2.3 Definitions ... 2-3
2.4 Polaris ... 2-4
2.4.1 Scheduled and Overall Availability ... 2-4
2.4.1.1 Explanation of Significant Availability Losses ... 2-4
2.4.2 System Mean Time to Interrupt (MTTI) and System Mean Time to Failure (MTTF) ... 2-6
2.4.2.1 MTTI and MTTF Summary ... 2-6
2.4.3 Resource Utilization ... 2-6
2.4.3.1 System Utilization ... 2-6
2.4.3.2 Capability Utilization ... 2-8
2.5 Theta ... 2-11
2.5.1 Scheduled and Overall Availability ... 2-11
2.5.1.1 Explanation of Significant Availability Losses ... 2-11
2.5.2 System Mean Time to Interrupt (MTTI) and System Mean Time to Failure (MTTF) ... 2-13
2.5.2.1 MTTI and MTTF Summary ... 2-13
2.5.3 Resource Utilization ... 2-14
2.5.3.1 System Utilization ... 2-14
2.5.3.2 Capability Utilization ... 2-16
2.6 Storage ... 2-19
2.6.1 Theta-fs0 Lustre File System ... 2-19
2.6.1.1 Scheduled and Overall Availability ... 2-19

Contents (Cont.)

Explanation of Significant Availability Losses 2-19
2.6.1.2 MTTI and MTTF 2-21
MTTI and MTTF Summary 2-21
2.6.2 Grand and Eagle Lustre File Systems 2-21
2.6.2.1 Grand Scheduled and Overall Availability 2-21
Grand – Explanation of Significant Availability Losses 2-22
Grand – MTTI and MTTF Summary 2-23
2.6.2.2 Eagle – Scheduled and Overall Availability 2-23
Eagle – Explanation of Significant Availability Losses 2-24
Eagle – MTTI and MTTF Summary 2-25
2.6.3 Tape Storage 2-25
2.6.3.1 Scheduled and Overall Availability 2-26
Explanation of Significant Availability Losses 2-26
2.6.3.2 MTTI and MTTF 2-28
MTTI and MTTF Summary 2-28
2.7 Center-Wide Operational Highlights 2-28
2.7.1 Polaris 2-28
2.7.2 AI Testbed 2-29
2.7.3 Sunspot / Aurora 2-29
2.7.4 Storage Changes 2-30
Conclusion 2-30

Section 3. Allocation of Resources 3-1

ALCF Response 3-1
3.1 Usage of the INCITE and ALCC Hours 3-1
3.2 Facility Director’s Discretionary Reserve Time 3-7
3.3 User Allocation Utilization 3-10
3.4 Other Large-Scale Managed Resources 3-10
Conclusion 3-11

Section 4. Operational Innovation 4-1

ALCF Response 4-1
4.1 Operational Innovation – Technical 4-1
4.1.1 Working to Enable On-Demand Workloads 4-1
4.1.2 Increased Availability of Logs Using Kafka 4-1
4.1.3 Making Filesystems a Scheduler-Aware Resource 4-2
4.1.4 HIP on Aurora 4-2
4.1.5 Benchmarking Resources in AI Testbed 4-3
4.2 Operational Innovation – Management/Workforce 4-3
4.2.1 Minimizing Cost and Effort around Managing SSL Certificates 4-3
4.2.2 ALCC and INCITE Project Introduction and Welcome Calls 4-3
4.2.3 AI-Related Training 4-4
4.3 Postdoctoral Fellows 4-4

Contents (Cont.)

Section 5. Risk Management 5-1

ALCF Response 5-1
5.1 Risk Management Process Overview 5-1
  5.1.1 Risk Review Board 5-3
  5.1.2 Risk Management in Day-to-Day Operations 5-3
  5.1.3 Continuation of the ALCF-3 Project 5-4
5.2 Major Risks Tracked 5-4
5.3 Risks Encountered in the Review Year and Their Mitigations 5-5
  5.3.1 Interruptions to Facility that Provides Cooling 5-5
  5.3.2 Interruptions to the ALCF Network 5-6
  5.3.3 Signing Delays of Some INCITE/ALCC User Agreements 5-7
  5.3.4 System Stability Issues 5-8
  5.3.5 Diagnostic Suite and Utilities Fail to Detect Hardware Problems 5-8
5.4 Retired Risks 5-9
5.5 New and Recharacterized Risks 5-9
5.6 Top Operating Risks Monitored Closely for the Next Year 5-9
Conclusion 5-10

Section 6. Environment, Safety, and Health 6-1

ALCF Response 6-1
6.1 Work Planning 6-1
6.2 COVID-19 Site Reentry 6-2
6.3 Continued Operations 6-2
6.4 Aurora Deliveries 6-2
6.5 Incident Reports 6-3

Section 7. Security 7-1

ALCF Response 7-1
7.1 Continual Improvement in Cyber Security Practices 7-1
7.2 Cyber Security Plan 7-4
7.3 Foreign Visitor Access and Export Controls 7-4

Section 8. Mission Impact, Strategic Planning, and Strategic Engagements 8-1

ALCF Response 8-1
8.1 Science Highlights and Accomplishments 8-1
  Science Highlights 8-3
  8.1.1 Computer Simulations Aided Design of Environmentally Benign Electrolyte for Lithium Battery 8-3
  8.1.2 Metastable Phase Diagrams for Materials 8-5
  8.1.3 Next-Generation Nonwovens Manufacturing: A Model-Driven Simulation and Machine Learning Approach 8-7
  8.1.4 Extreme-Scale Simulations for Advanced Seismic Ground Motion and Hazard Modeling 8-9
  8.1.5 Binding Energies from Effective Field Theory with Quantified Uncertainties 8-11

Contents (Cont.)

8.2 Research Activities / Vendor Engagements for Future Operations ... 8-13
  8.2.1 Research Activity - Joint Laboratory for System Evaluation (JLSE) ... 8-13
  8.2.2 Research Activity - ALCF AI Testbed ... 8-15
  8.2.3 Vendor Engagement – Codeplay Software Ltd. (Codeplay) ... 8-17
  8.2.4 Vendor Engagement – Altair and OpenPBS ... 8-17
  8.2.5 Vendor Engagement – DOE Advanced Computing Ecosystem RFI ... 8-17
8.3 DOE Program Engagements / Requirements Gathering ... 8-18
  8.3.1 Engagement Highlights ... 8-18
    Supercomputing 2022 ... 8-18
    Engagement in Standards and Community Groups ... 8-19
    Performance, Portability, and Productivity in HPC (P3HPC) Forum ... 8-19
  8.3.2 Summary of Engagements with the Exascale Computing Project ... 8-19
    ECP-Funded Positions in ALCF ... 8-20
    Continuous Integration (CI) Pipeline ... 8-20
    Communication between the ALCF and the ECP Resource Allocation Council ... 8-20
Conclusion ... 8-20

Appendix A – Calculations ... A-1

A.1 Scheduled Availability ... A-1
A.2 Overall Availability ... A-1
A.3 System Mean Time to Interrupt (MTTI) ... A-2
A.4 System Mean Time to Failure (MTTF) ... A-2
A.5 Total System Utilization ... A-2
A.6 Capability ... A-2
A.7 Theta Nodes ... A-3

Appendix B – ALCF Director's Discretionary Projects ... B-1

Appendix C – ALCF CY2022 Science Highlights ... C-1

Appendix D – Acronyms and Abbreviations ... D-1

Figures

2.1 Polaris Weekly Availability for CY 2022 ... 2-5

2.2 Polaris System Utilization over Time by Program ... 2-7

2.3 Polaris INCITE Overall Capability ... 2-9

2.4 Polaris Capability Node-Hours by Program ... 2-10

2.5 Polaris Job Usage by Size ... 2-10

2.6 Theta (with expansion) Weekly Availability for CY 2022 ... 2-12

2.7 System Utilization over Time by Program ... 2-15

Figures (Cont.)

2.8 Theta INCITE Overall Capability 2-17

2.9 Theta Capability Node-Hours by Program 2-18

2.10 Theta Job Usage by Size 2-18

2.11 Theta-fs0 Weekly Availability for CY 2022 2-20

2.12 Grand Weekly Availability for CY 2022 2-22

2.13 Eagle Weekly Availability for CY 2022 2-24

2.14 HPSS Weekly Availability for CY 2022 2-27

3.1 Theta (with expansion) INCITE 2022 Allocation Usage (Note: Projects are randomly ordered.) 3-3

3.2 Polaris INCITE 2022 Allocation Usage (Note: Projects are randomly ordered.) 3-4

3.3 Theta (with expansion) ALCC 2021–2022 Allocation Usage (Note: Projects are randomly ordered.) 3-5

3.4 Theta (with expansion) ALCC 2022–2023 Allocation Usage as of December 31, 2022 (Note: Projects are randomly ordered.) 3-5

3.5 Polaris ALCC 2022–2023 Allocation Usage as of December 31, 2022 (Note: Projects are randomly ordered.) 3-6

3.6 Theta (with expansion) CY 2022 DD Allocations by Domain 3-9

3.7 Polaris CY 2022 DD Allocations by Domain 3-9

5.1 Illustration of the Iterative Risk Management Process from ALCF's Risk Management Plan 5-2

8.1 Simulated cluster of water (red) and lithium ions (green) at electrode surfaces under various water concentrations. Battery performance is retained when water molecules remain isolated. 8-3

8.2 The final product of the machine learning algorithm: metastable phase diagrams for carbon. The colored regions indicated conditions at which carbon exists in certain metastable states (with similarly colored structures) that may yield useful material properties. 8-5

8.3 (top) Melt-blowing experiments at 3M and (bottom) Argonne CFD simulation results modeling the same process using conditions selected from machine learning. 8-7

8.4 Simulation model setup showing (a) observational and model fault details with arrows indicating orientations of principal stresses, (b) fault-plane model and epicenter (star), and (c) rock properties. 8-9

8.5 Calculated ground state energies in MeV using chiral interactions in comparison with experimental values. 8-11

Figures (Cont.)

8.6 Breakdown of some key activities by ALCF in CY 2022. The first pie chart (left) breaks down the 191 events by type, primarily derived by how the event identified itself. The second pie chart (right) breaks down the same events by the role of the staff member. 8-18

Tables

ES.1 Summary of the Target and Actual Data for the Previous Year (2022) Metrics ES-2
RTR.1 ALCF Response to Recommendation from the PY OAR RTR-2
1.1 All 2022 User Support Metrics and Results 1-1
1.2 User Survey Results in 2022 by Allocation Program 1-4
1.3 User Support Metrics 1-4
1.4 Problem Resolution Metrics 1-5
1.5 Ticket Categorization for 2021 and 2022 1-6
1.6 Publications Designed for Print 1-15
1.7 Primary Communication Channels 1-16
1.8 Target Audiences 1-16
2.1 Summary of Operational Performance of HPC Computational Systems 2-1
2.2 Summary of Operational Performance of HPC Storage Systems 2-2
2.3 Availability Results 2-4
2.4 MTTI and MTTF Results 2-6
2.5 System Utilization Results 2-7
2.6 Node-Hours Allocated and Used by Program 2-8
2.7 Capability Results 2-9
2.8 Availability Results 2-11
2.9 MTTI and MTTF Results 2-14
2.10 System Utilization Results 2-15
2.11 Node-Hours Allocated and Used by Program 2-16
2.12 Capability Results 2-17
2.13 Availability Results 2-19
2.14 MTTI and MTTF Results 2-21
2.15 Availability Results – Grand 2-22
2.16 MTTI and MTTF Results – Grand 2-23

Tables (Cont.)

2.17 Availability Results – Eagle 2-24
2.18 MTTI and MTTF Results – Eagle 2-25
2.19 Availability Results 2-26
2.20 MTTI and MTTF Results 2-28
3.1 INCITE 2022 Time Allocated and Used on Theta (with expansion) in CY 2022 3-3
3.2 INCITE 2022 Time Allocated and Used on Polaris in CY 2022 3-4
3.3 ALCC Time Allocated and Used on Theta (with expansion) in CY 2022 3-6
3.4 ALCC Time Allocated and Used on Polaris in CY 2022 3-6
3.5 DD Time Allocated and Used on Theta (with expansion) in CY 2022 3-8
3.6 DD Time Allocated and Used on Polaris in CY 2022 3-8
5.1 Major Risks Tracked for CY 2022 5-5
5.2 Top Operating Risks Monitored for CY 2023 5-9
6.1 ALCF Injury and Illness Cases FY 2018 thru FY 2022 6-3
8.1 Summary of Users' Peer-Reviewed Publications in CY 2022 8-2
8.2 Summary of Users' Peer-Reviewed Publications for 5-year Moving Window 8-2
8.3 Summary of ALCF Participation in SC22 8-19
A.1 Capability Definitions for Theta A-2
A.2 Total and Reportable Nodes for Theta A-3

Background

In 2004, the U.S. Department of Energy's (DOE's) Advanced Scientific Computing Research (ASCR) program established 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 large-scale 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 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 operate around the clock every day of the year. From an operational standpoint, the high level of services these centers provide and the exceptional science they produce justify their existence to the DOE Office of Science and the U.S. Congress.

This page intentionally left blank.

Executive Summary

This Operational Assessment Report describes how the Argonne Leadership Computing Facility (ALCF) met or exceeded every one of its goals for the calendar year (CY) 2022.

In CY 2022, ALCF operated Theta, an Intel-based Cray XC40 system (11.7 petaflops) augmented with 24 NVIDIA DGX A100-based nodes (3.9 petaflops), that supports diverse workloads, integrating data analytics with artificial intelligence (AI) training and learning in a single platform; and Polaris, a 44-petaflop AMD and NVIDIA-based HPE Apollo 6500 Gen10+ system that provides a powerful new platform to prepare applications and workloads for Aurora, Argonne National Laboratory's (Argonne's) upcoming Intel-Hewlett Packard Enterprise (HPE) exascale supercomputer.

Last year, Theta delivered a total of 20.4 million node-hours to 17 Innovative and Novel Computational Impact on Theory and Experiment (INCITE) projects and 7.4 million node-hours to ASCR Leadership Computing Challenge (ALCC) projects (18 awarded during the 2021–2022 ALCC year and 11 awarded during the 2022–2023 ALCC year), as well as provided substantial support to Director's Discretionary (DD) projects (5.3 million node-hours). As Table ES.1 shows, Theta performed exceptionally well in terms of overall availability (95.7 percent), scheduled availability (98.6 percent), and utilization (97.0 percent; Table 2.1).

Upon going live in August, Polaris began supporting several research teams via the U.S. Department of Energy's (DOE's) Exascale Computing Project (ECP) and the ALCF's Aurora Early Science Program (ESP). It also supported ALCC and INCITE projects, delivering a total of 589.6K node-hours to 10 INCITE projects and 92.6K node-hours to 6 ALCC projects awarded during the 2022–2023 ALCC year (Polaris was not available for the 2021–2022 ALCC year), and also supported DD projects (543.6K node-hours). As Table ES.1 shows, Polaris performed extremely well in terms of overall availability (94.7 percent), scheduled availability (98.2 percent), and utilization (67.3 percent; Table 2.1).

In CY 2022, ALCF supported more than 1,500 users, including those who came for training purposes, and as of the submission date of this document, ALCF's users have published 213 papers in peer-reviewed journals and technical proceedings. Other 2022 highlights include:

- An Argonne-led team won the ACM Gordon Bell Special Prize for HPC-based COVID-19 Research at the 2022 International Conference for High Performance Computing, Networking, Storage, and Analysis (SC22), for work that used ALCF resources to create the first genome-scale language models (GenSLMs) for understanding the evolution of SARS-CoV-2.
- Two systems in ALCF's AI Testbed were opened to the entire research community.
- ALCF moved all user documentation to GitHub to enhance the overall user experience.
- ALCF hosted the most successful AI Training program yet to help train over 200 current and potential users.

In 2022, national laboratories across the country officially entered the exascale era, with Aurora at Argonne on the horizon. ALCF researchers will continue to engage in strategic activities that push the boundaries of what's possible in computational science and engineering and allow ALCF to deliver science on day one.

Table ES.1 Summary of the Target and Actual Data for the Previous Year (2022) Metrics

Area	Metric	2022 Target	2022 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%	38.4%
	% User Problems Addressed within Three Working Days	80.0%	93.3%
Business Results	Polaris Overall Availability	90.0%	94.7%
	Polaris Scheduled Availability	90.0%	98.2%
	Theta (with expansion) Overall Availability	90.0%	95.7%
	Theta (with expansion) Scheduled Availability	90.0%	98.6%
	% of INCITE node hours from jobs run on 20.0% or more of Polaris (99–532 nodes)	20.0%	47.7%
	% of INCITE node hours from jobs run on 60.0% or more of Polaris (297–532 nodes)	N/A [b]	22.3%
	% of INCITE node hours from jobs run on 20.0% or more of Theta (800–4008 nodes)	20.0%	85.8%
	% of INCITE node hours from jobs run on 60.0% or more of Theta (2400–4008 nodes)	N/A	27.2%
	Theta-fs0 Overall Availability	90.0%	96.3%
	Theta-fs0 Scheduled Availability	90.0%	99.3%
	Grand Overall Availability	90.0%	96.6%
	Grand Scheduled Availability	90.0%	99.5%
	Eagle Overall Availability	90.0%	96.6%
	Eagle Scheduled Availability	90.0%	99.5%
	HPSS Overall Availability [a]	90.0%	96.7%
	HPSS Scheduled Availability	90.0%	99.7%

[a] HPSS = high-performance storage system. (Appendix D contains a list of the acronyms and abbreviations used in this report.)

[b] N/A = not applicable.

Response to Recommendation from the PY OAR

The Facility should list recommendations from the previous year's OAR and the Facility's responses to them. If prior year recommendations were directed to DOE/ASCR, that point should be clearly stated. The Facility should include the responses from DOE/ASCR in the report.

The Facility may also comment on any plans or actions that may have changed since these original responses.

NOTE: This data is for informational purposes, and is not formally a part of the review. It will allow the reviewers to place certain actions in context.

Table RTR.1 (on the following page) summarizes ALCF's response to a recommendation from the previous year (PY) OAR.

Table RTR.1 ALCF Response to Recommendation from the PY OAR

Recommendation	Facility Response	DOE/ASCR Response
The report omits a response to the recommendation that was delivered during the onsite review in March 2021 under charge question 7(a), stating "In light of the rapidly evolving technology landscape, application types, and scientific workflows, analyze how cybersecurity requirements will evolve, how security monitoring and other activities at ALCF compare with peer facilities, the risk of staff loss and succession planning, and consider cyber security staffing levels." While it is true that ALCF continues to make notable improvements in security, threats also continue to intensify, as evidenced by the breadth and severity of the Log4Shell exploit in December 2021. We urge ALCF to conduct and document an analysis of staffing levels in this area.	ALCF neglected to include its previous year's (PY) response to recommendations in the 2021 OAR review and will correct this omission in 2022. Regarding the 2021 recommendation, ALCF disagrees with the underlying premise that ALCF staffing levels must be justified without any reasoning beyond the reviewer's concerns about the evolving technology landscape. Moreover, no details on staffing are requested, or provided, as part of the OAR charge question concerning security; therefore, it is difficult for ALCF to respond. ALCF has bench depth in cybersecurity beyond that of a single dedicated officer. The facility has two additional staff members who serve as backup cybersecurity program representatives, similar to peer facilities. In addition, ALCF's Operations team has established strong cybersecurity practices analogous to the excellent physical safety and security practices ALCF already maintains. ALCF, at ASCR's request, benchmarked its cybersecurity staffing numbers against other peer facilities with similar user populations (OLCF, TACC) in July 2022 and found them comparable. ALCF has also discussed its cybersecurity strategy with Computing, Environment, and Life Sciences Information Technology (CELS IT, ALCF's home directorate) and the Argonne Cybersecurity Office, and neither reported any significant concerns (please see the Division Site Assist Visit (DSAV) report in this year's review). ALCF staff are also actively engaged in the ASCR-led Secure ASCR Facilities (SECAF) working group activities. As previously mentioned, the cybersecurity function is embedded within all staff roles at ALCF. Our web developers use best practices such as role-based access control and performing all input validation on the backend. Our system administrators install new software and updates on test systems first, whenever possible, and review anything from an untrusted source before installation. Based on risk ratings, common vulnerabilities and exposures are addressed for each system following the best practices and mandates of the U.S. Department of Homeland Security. In addition, the ALCF is also forward-looking as it works to hire for a cybersecurity software developer position to bring additional cybersecurity controls to tools and projects that are developed in-house.	While ALCF neglected to include its response to the previous year's recommendation to the review copy of the OAR, they submitted this response to ASCR. ALCF also presented a plan to the ASCR PM in summer of 2022 about how they plan to follow up on this recommendation (beyond just including responses in the version available to the reviewers) and other items deemed actionable by the ASCR PM. In July 2022, ALCF verbally informed ASCR that comparable Facilities, such as OLCF and TACC, have similar levels of cybersecurity staffing and protocols.

Section 1. User Support Results

Are the processes for supporting users, resolving users' problems, and conducting outreach to the user population effective?

ALCF Response

ALCF has processes in place to effectively support its customers, to resolve user problems, and to conduct outreach. The 2022 annual user survey measured overall satisfaction, user support, and problem resolution, and serves both to mark progress and to identify areas for improvement (Table 1.1). User satisfaction with ALCF services remains 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) 2022.

Table 1.1 All 2022 User Support Metrics and Results [a]

		2021 Target	2021 Actual	2022 Target	2022 Actual
Number Surveyed		N/A [c]	1,174	N/A [c]	1,446
Number of Respondents (Response Rate)		25.0%	483 (42.3%)	25.0%	555 (38.4%)
Overall Satisfaction	Mean	3.5	4.5	3.5	4.5
	Variance	N/A	0.5	N/A	0.6
	Standard Deviation	N/A	0.7	N/A	0.8
Problem Resolution	Mean	3.5	4.6	3.5	4.5
	Variance	N/A	0.5	N/A	0.5
	Standard Deviation	N/A	0.7	N/A	0.7
User Support	Mean	3.5	4.6	3.5	4.5
	Variance	N/A	0.5	N/A	0.5
	Standard Deviation	N/A	0.7	N/A	0.7
		2021 Target	**2021 Actual**	**2022 Target**	**2022 Actual**
% of Problems Addressed Within Three Working Days [b]		80.0%	95.0%	80.0%	93.3%

[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 the period in question. This definition of a user provides the basis for all survey results.

[b] The statistical population represented in this metric includes problem tickets submitted from all users. Note that this is a larger population than those who qualify under the September 2015 definition of a facility user mentioned in note *a* above.

[c] N/A = not applicable.

Survey Approach

In 2017, the ALCF worked with a consultancy to revise and shorten its annual user survey: omitting the workshop-related questions, requiring responses only for those questions that comprise the DOE metrics for the Operational Assessment Report (OAR), and making every other question optional. The facility polls workshop attendees separately.

The 2022 user survey closely resembled the 2021 survey, with a few modifications to the operations, infrastructure, and science and technical support areas listed in various questions. Four new questions were added: one that captured the use or planned use of AI on ALCF systems (including the AI testbed), and three questions around training events at the ALCF. No questions were removed.

The 2022 survey was administered through a contract with Statistical Consulting Services at the Department of Statistics and Actuarial Science at Northern Illinois University and consisted of 6 required questions and 24 optional questions. The survey and associated e-mail campaign ran from November 16, 2022, through December 31, 2022. Each reminder e-mail included a user-specific link to the online survey. Most respondents were able to complete the survey in 10 minutes or less.

Likert Scale and Numeric Mapping

Almost all questions in the survey used a six-point Likert 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.

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.

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

Everyone who met the definition of a facility user during the fiscal year (FY) 2022—that is, users who would be counted under the Facility's annual user statistics submission to the Office of Science—was asked to complete a user survey, 1,538 individuals in total. Of those individuals, 16 did not receive the email due to undeliverable messages and 76 chose to opt-out of the survey. Of the 1,446 remaining users, 555 responded, for a 38.4 percent response rate. The ALCF surpassed all targets for the survey metrics.

Table 1.2 shows user survey results grouped by allocation program. While Innovative and Novel Computational Impact on Theory and Experiment (INCITE) and ASCR Leadership Computing Challenge (ALCC) users reported higher average Overall Satisfaction than Director's Discretionary (DD) users, the variations are very minor. Other metrics are similar, in that the variations are statistically insignificant.

Table 1.2 User Survey Results in 2022 by Allocation Program

Metrics by Program		INCITE	ALCC	INCITE + ALCC	DD	All
Number Surveyed		206	100	306	1,140	1,446
Number of Respondents		96	51	147	408	555
Response Rate		46.6%	51.0%	48.0%	35.8%	38.4%
Overall Satisfaction	Mean	4.6	4.5	4.6	4.5	4.5
	Variance	0.5	0.5	0.5	0.6	0.6
	Standard Deviation	0.7	0.7	0.7	0.8	0.8
Problem Resolution	Mean	4.6	4.6	4.6	4.5	4.5
	Variance	0.5	0.4	0.5	0.5	0.5
	Standard Deviation	0.7	0.6	0.7	0.7	0.7
User Support	Mean	4.6	4.5	4.6	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	4.5	4.5	4.6	4.5	4.5
	Variance	0.5	0.5	0.5	0.6	0.5
	Standard Deviation	0.7	0.7	0.7	0.7	0.7

As Table 1.3 shows, in 2022, the ALCF again exceeded the targets for overall satisfaction and user support.

Table 1.3 User Support Metrics

Survey Area	2021 Target	2021 Actual	2022 Target	2022 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	4.6/5.0	3.5/5.0	4.5/5.0

1.2 Problem Resolution Metrics

Table 1.4 shows the target set for the percentage of user problem tickets or “trouble tickets” addressed in three days or less, which ALCF exceeded. A trouble ticket, which encompasses incidents, problems, and service requests, is considered “addressed” once (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 Problem Resolution Metrics

	2021 Target	2021 Actual	2022 Target	2022 Actual
% of Problems Addressed Within Three Working Days [a]	80.0%	95.0%	80.0%	93.3%
Avg. of Problem Resolution Ratings	3.5/5.0	4.6/5.0	3.5/5.0	4.5/5.0

[a] The statistical population represented in this metric includes all users that submitted a ticket. Note that this is a larger population than those who qualify under the September 2015 definition of a facility user (see note *a* in Table 1.1).

1.3 User Support and User Engagement

1.3.1 Tier 1 Support

1.3.1.1 Phone, Slack, and E-mail Support

In 2022, ALCF addressed and categorized 5,801 user support tickets, a 16 percent increase from the previous year. The biggest increases were in the Accounts and Allocations categories (Table 1.5). ALCF rolled out multiple new systems—SambaNova, Cerebras, and Polaris—in 2022 (see Section 2.2 for Polaris; and Section 8.2.2 for SambaNova and Cerebras), resulting in additional allocation and account request activity. Multiple training events for these new systems contributed to an increase in account requests.

In addition, ALCF's "Introduction to AI-driven Science on Supercomputers" training event resulted in more than 400 account requests that were not tracked via traditional support tickets and are not included in the table below. Instead, everyone who requested an account received one, provided they included their university email address in their application. In this process, a ticket is only generated when a project PI/proxy explicitly approves a request. In the absence of an explicit approval for each request from the training leads (who are the project PI/proxies), no tickets were generated.

ALCF also provided access to early users and vendor staff to Sunspot that were not tracked via support tickets (see Section 2.7.3 for Sunspot). These two exclusions would have added at least an additional 600 account/access tickets that are not reflected in the table below. ALCF provided real-time account and access support via Slack workspaces for many of its CY 2022 training events.

Table 1.5 Ticket Categorization for 2021 and 2022

Category	2021	2022
Access	808 (16%)	915 (15%)
Accounts	2,084 (40%)	2,490 (42%)
Allocations	856 (16%)	1080(18%)
Applications Software	172 (3%)	201 (3%)
Automated E-mail Responses	205 (4%)	88 (1%)
Compilers	37 (1%)	35 (1%)
Data Transfer	54 (1%)	49 (1%)
Debugging and Debuggers	9 (0%)	7 (0%)
File System	162 (4%)	66 (1%)
HPSS [a] and Quota Management	102 (2%)	156 (2%)
Libraries	36 (1%)	20 (1%)
Miscellaneous	148 (3%)	119 (2%)
Network	9 (0%)	8 (0%)
Performance and Performance Tools	15 (0%)	9 (0%)
Reports	181 (4%)	146 (2%)
Scheduling	154 (3%)	337 (5%)
System [b]	N/A [c]	73 (1%)
Visualization	2 (0%)	2 (0%)
TOTAL TICKETS	5,034 (100%)	5,801 (100%)

[a] HPSS = high-performance storage system.

[b] The System category is a new category this year.

[c] N/A = not applicable.

1.3.1.2 User Account Management Software (Userbase3) Improvement

Argonne is a controlled-access facility, and anyone entering the site or accessing Argonne resources remotely must be authorized. Users can access ALCF resources only with an active ALCF account. To apply for an ALCF account, the user fills out a secure webform in the ALCF Account and Project Management system, Userbase 3 (UB3).

Significant Improvements in Automated Testing for Userbase3

A virtual machine was set up to do nightly runs of automated tests. This ensures that feature changes and fixed software issues do not affect other functionality, providing for a better released software product and for a way to run tests consistently instead of relying on machine-to-machine timing, variation, etc. ALCF increased the number of automated tests by over 40 percent in 2022, with a failure rate of under 10 percent. Testing improvements were made by doing more event-driven tests instead of event waiting and by writing tests using generic data.

Improvements in Process for Access Control and NDA Projects in Userbase3

An Access Control Screen was added that allowed ALCF administrators to add access control to projects, Unix groups, and systems using a graphical user interface (GUI) instead of using backend calls. A new software feature was also added to help with the vetting process when adding a new user to Nondisclosure Agreement (NDA) projects.

Support for AI Testbeds and Polaris in Userbase3

Allocation requests for Polaris and the AI Test Bed machines (SambaNova and Cerebras) were added as options to the Allocation Request webpage. Userbase3 collects Director Discretionary allocation requests and sends them to the Allocation Request Management system for approval.

Auto Subscribing Users to System Mailing Lists

In the past, when users' accounts were activated at ALCF, they would receive an invitation to join the system mailing lists. This process has been improved: instead of being invited to subscribe, users are now auto subscribed to a given system's mailing lists as soon as their accounts are activated for that new system. This modification ensures timely communication about ALCF systems to all users.

1.3.1.3 Performance Portability Study

The performance of applications ported across systems with different architectures is often not well documented nor captured. Developers do not have enough information to improve software stacks. ALCF personnel worked closely with Argonne Computational Science (CPS) Division staff members to evaluate progress toward achieving performance portability on AMR-Wind, HACC, SW4, GAMESS RI-MP2, XSBench and TestSNAP, with performance portability assessed across the AMD MI100, Intel Gen9, and NVIDIA A100 graphics processing unit (GPU) platforms. The staff used the roofline performance model to compute performance efficiency and evaluate performance portability across the three platforms. Results were first presented and published at the 2021 International Workshop on Performance, Portability and Productivity in HPC (P3HPC) held at SC21. This work has continued in CY 2022 with the intent of publishing an additional paper comparing performance of a similar set of applications on the latest generation of GPUs.

1.3.1.4 Gaussian Process Modeling at HPC Scales

Deep learning can often be challenging when dimensions/parameters are very large. The "curse of dimensionality" is one such limitation that has prevented Gaussian process (GP) modeling methods from being applied in big data contexts. However, new, fast, and accurate approximation methods with clear mathematical foundations and good explainability properties have broken through this limitation. GPyTorch is a package that implements such methods. By deploying the GPyTorch library on ALCF machines, ALCF provides users with access to this kind of modeling on a GPU-accelerated software platform for GPU-based AI work, encouraging users to use this novel capability in their own applications.

1.3.2 Application Support for Individual Projects

Enabling analysis of WRF outputs for ALCC project

ALCF staff worked with the ALCC project MesoConvSyss to enable more effective post-processing of large volumes of output of Weather Research and Forecasting (WRF) simulations (e.g., a single ~300 GB file generated from 16,384 separate WRF files). This required migrating to an MPI-parallel version of join WRF analysis code (with some edits) to accommodate memory requirements and use of multiple compute nodes. This assistance enabled the MesoConvSyss team to quickly process its massive simulation output.

Initial porting of LAMMPS code to GPUs as part of DD project for future ALCC Submission

ALCF staff worked with the DD project XMultiImage to develop an initial GPU port of the LAMMPS MC/MD code. While implementation is still early, using ThetaGPU the project is already seeing an overall application speedup of ~32x on an A100 GPU compared to a single central processing unit (CPU). More work is needed to optimize data transfers, offload additional code, and generate proper benchmark comparisons. The GPU-enabled code will be used in future ALCC submissions targeting Polaris and Aurora.

1.3.3 User Engagement

1.3.3.1 User Advisory Council

ALCF's User Advisory Council (UAC) provides guidance on proposed policy changes and services and gives feedback on the experiences and needs of the user community in general. Members are appointed by the ALCF Director and have expert knowledge of the tasks and requirements of specific applications or domain areas. In CY 2022, council members Sibendu Som and Aiichiro Nakano were reappointed, and five others rotated off: C.S. Chang, Jacqui Cole, William Detmold, Arvind Ramanathan, and Tom LeCompte. Meetings have been suspended pending a planned reboot of the council in 2023.

1.3.3.2 Training Activities

ALCF offers workshops and webinars on various tools, systems, and frameworks. These hands-on training programs are designed to help PIs, their project members, and future users take advantage of leadership-class computers available at ALCF and enhance the performance and productivity of their research. ALCF also collaborates with peer institutions and vendor partners to offer training that strengthens community competencies and promotes best practices. In CY 2022, ALCF conducted over 30 training activities, reaching more than 500 participants and 25 teams in events where participation was tracked. Below is a list of ALCF 2022 training activities:

ALCF GPU Hackathon (virtual)

ALCF and NVIDIA hosted a multi-day GPU hackathon involving 13 teams of developers who worked with assigned mentors to accelerate their codes on ThetaGPU using portable programming models such as OpenMP or an AI framework. (Dates: April 20, and April 27–29, 2022)

Computational Performance Workshop (virtual)

This annual workshop helps science teams achieve computational readiness for INCITE and other ALCF allocation programs. In 2022, participants worked with ALCF and vendor staff to debug/benchmark codes on ALCF resources and expand their data science skills. The attendees came from a wide variety of institutions: ALCF/CPS (47), other Argonne (7), industry (14), DOE labs (16), and universities (30). Of the non-Argonne attendees who completed the feedback survey, nearly half said they plan to apply for an INCITE award. (Dates: May 24–26, 2022)

2022 INCITE Proposal Writing Webinars (virtual)

The INCITE program, ALCF, and the Oak Ridge Leadership Computing Facility (OLCF) jointly hosted two webinars on effective strategies for writing an INCITE proposal. (Dates: May 3, 2022, and June 3, 2022)

2022 NVIDIA ALCF GPU Hackathon (virtual)

ALCF and NVIDIA hosted a multi-day GPU hackathon involving 12 teams of 3–6 developers each, and their assigned mentors. It was the first such event to provide access to Polaris. Attendees scored Argonne 4.5/5.0 in providing satisfactory support at the hackathon. (Dates: July 19 and July 26–28, 2022)

Argonne Training Program on Extreme-Scale Computing (ATPESC)

ATPESC, founded by Argonne and now part of the Exascale Computing Project, is a two-week intensive training program on the key skills and tools needed to use supercomputers for science. The program features talks given by leading computer scientists and HPC experts; and hands-on training using DOE leadership-class systems at ALCF, OLCF, and the National Energy Research Scientific Computing Center (NERSC). In 2022, ATPESC attracted 79 attendees from 58 different institutions worldwide. Video recordings of ATPESC sessions are available on Argonne's YouTube training channel. (Dates: July 31–August 12, 2022)

Simulation, Data, and Learning (SDL) Workshop (virtual)

The annual SDL interactive workshop is aimed at researchers who are planning to apply for a major allocation award in the near term. Participants learn to scale data-centric science on ALCF systems, set up workflows, use containers, and test and debug codes. This year's workshop focused on Polaris, ThetaGPU, and resources in ALCF's AI Testbed. Ninety attendees joined the first day: ALCF/CPS (29), other Argonne staff (5), other institutions (56). (Dates: October 4–6, 2022)

Intro to AI-driven Science on Supercomputers: A Student Training Series (virtual)

The AI-driven Science on Supercomputers 8-week webinar series is aimed at undergraduate and graduate students enrolled at U.S. universities and community colleges and designed to attract a new generation of AI users by having a low entry barrier; that is, attendees need to have only a basic experience with the Python programming language as the pre-requisite.

The 2022 training series incorporated feedback from participants from the previous ALCF AI for Science Training Series (October 2021–February 2022) and input from Argonne's Educational Programs and Outreach office. ALCF computer scientists led the weekly sessions and hands-on exercises along with talks by Argonne scientists who use AI in their work. The 2022 program welcomed more than 200 attendees from 90 universities, including undergraduates, graduate

students, postdocs, and faculty. Of those attendees, 52 received a Certificate of Completion for the course, completing all eight homework assignments from each session.
(Dates: September 2022–November 2022)

Aurora Early Science Program (ESP) Dungeons, Hackathons, and Workshops (virtual)

In 2022, the Intel Center of Excellence (COE), in collaboration with ALCF's Early Science Program (ESP), held multiday events where select ESP and ECP project teams worked on developing, porting, and profiling their codes with help from Intel and Argonne experts. The 2022 program included the following activities:

- March 16–18, 2022: Hackathon 16 (Fusion Energy)
- March 16–18, 2022: Dungeon 3 (projects CANDLE, Uintah, Fusion Energy, PHASTA Sim (generating data for the ML [Machine Learning] project) & PHASTA ML)
- May 9–10, 2022: Hackathon 17 (HACC)
- June 14–16, 2022: Workshop #4 (Aurora)
- September 1–2, 2022: Hackathon 18 (LQCD)
- October 26–27, 2022: Hackathon 19 (HARVEY)
- December 13–15, 2022: Dungeon 4

ALCF Webinars (virtual)

The 2022 ALCF webinar program consisted of two tracks: *ALCF Developer Sessions* and *Aurora Early Adopter Series*. The ALCF Developer Sessions were focused on those writing code for Aurora. The Aurora Early Adopter Series was focused on public discussions related to Aurora. All talks are posted to ALCF's YouTube channel, and the associated training materials can be found on the ALCF Events website. ALCF also participates in useful community events, the IDEAS productivity project webinar series, and Intel webinars. The 2022 webinar program was as follows:

- January 26, 2022: Getting Started on ThetaGPU
- February 23, 2022: NVIDIA Performance Tools for A100 GPU Systems
- March 30, 2022: Data Parallel Python: Bringing oneAPI to Python
- April 27, 2022: Polaris Overview
- May 25, 2022: The LLVM/OpenMP Ecosystem – Optimizations, Features and Outlook
- June 30, 2022: Profiling Deep Learning Applications with NVIDIA Nsight
- July 27, 2022: An Introduction to HDF5 for HPC Data Models, Analysis, and Performance
- August 31, 2022: HDF5 Workshop
- September 28, 2022: Getting Started on Polaris
- October 26, 2022: Intro to Intel Extensions of Scikit – learn to Accelerate Machine Learning Frameworks

- November 30, 2022: Debugging GPU-Accelerated Applications with NVIDIA Developer Tools

AI/Machine Learning (ML) Workshops/Training

ALCF invited researchers from across the laboratory to learn about the resources available in ALCF's AI Testbed and how they provide next-generation capabilities for science.

- April 26, 2022: ALCF DeepHyper Automated Machine Learning Workshop (public event)
- June 14–15, 2022: Habana – ALCF AI Testbed Training Workshop (NDA event)
- July 15, 2022: ALCF PythonFOAM Workshop (public event)
- July 19, 2022: SambaNova – ALCF AI Testbed Training Workshop (public event)
- July 25–August 1, 2022: Cerebras HPC SDK – ALCF AI Testbed Training Workshop (NDA event)
- August 9–10, 2022: Cerebras – ALCF AI Testbed Training Workshop (public event)

1.3.3.3 Community Outreach

In CY 2022, ALCF hosted 8 community outreach events with more than 350 users and reached 2,000+ more students in their classrooms. These activities ranged from giving tours to industry groups and DOE leadership to participating in science, technology, engineering and math (STEM) efforts and classroom visits directed at K-12 students. ALCF staff ran several summer coding camps and participated in annual computer science education events such as the Hour of Code. ALCF staff contribute to a wide range of activities aimed at sparking students' interest in scientific computing and promoting career possibilities in STEM fields. Additionally, the ALCF's annual summer student program gives college students the opportunity to work side-by-side with staff members on real-world research projects and to utilize some of the world's most powerful supercomputers, working in areas like computational science, system administration, and data science.

STEM Activities

Summer 2022 Research Internships

ALCF hosted 27 student interns through various programs including DOE's Science Undergraduate Laboratory Internships (SULI) program, Argonne's Research Aide program, the National Consortium for Graduate Degrees for Minorities in Engineering and Science (The National GEM Consortium), and the Professional Career Internship (PCI) program. Interns worked on mentored research projects in the field of scientific computing. These junior researchers used online collaboration tools to meet and conduct hands-on activities throughout the summer and presented their findings to the ALCF community in a seminar series at the end of their time at Argonne. The student projects included using AI to analyze bird songs, visualizing large scientific datasets, benchmarking graph neural networks for science on AI accelerators, and advancing high-energy physics research.

2022 Introduce a Girl to Engineering Day

Argonne hosted 140 eighth-grade girls for the 20th annual Introduce a Girl to Engineering Day (IGED). Seventy-three Argonne mentors, including ALCF staff members, volunteered for an engaging day of presentations and hands-on activities focused on STEM careers.
(Date: February 26, 2022)

2022 Science Careers in Search of Women

Argonne's 2022 Science Careers in Search of Women conference was held in-person for the first time in three years. This annual lab-sponsored event offered female high school students an opportunity to explore the world of STEM research through various interactions with Argonne's women scientists and engineers, including ALCF staff, with activities such as career panel discussions and poster presentations. A total of 110 girls participated in this event.
(Date: April 29, 2022)

2022 CodeGirls @ Argonne Camp (virtual)

Argonne held its sixth annual CodeGirls@Argonne summer camp, a weeklong STEM course for 6th- and 7th-grade girls taught by Argonne Learning Center and ALCF staff. The girls learned Python coding fundamentals (prior programming experience was not a pre-requisite), experimented with robotics, and met women scientists who use code to solve problems. The group also virtually toured the ALCF machine room and learned about the future Aurora supercomputer. (Dates: June 27–July 1, 2022)

Argonne-NIU AI Camp

Argonne and Northern Illinois University (NIU) hosted the second AI-focused summer camp for regional middle school and high school students recruited through NIU's Upward Bound program (prior programming experience was not a pre-requisite). During the monthlong camp, students were introduced to fundamental concepts of AI and machine learning and participated in hands-on activities for exploring sensor-collected datasets using software analysis tools. The camp was taught by scientific and educational outreach staff from Argonne and STEM educators from NIU and the University of Illinois Chicago. (Date: July 2022)

2022 Coding for Science Camp

Coding for Science Camp was a five-day enrichment experience for high school freshmen and sophomores who are new to coding. The camp curriculum promotes problem-solving and teamwork skills through hands-on coding activities, such as coding with Python and programming a robot, and interactions with Argonne staff members working in HPC and visualization. This camp, a joint initiative of Argonne's Educational Programs Office and ALCF, hosted 22 students last summer. (Dates: July 18–22, 2022)

2022 Big Data Camp

Argonne's fifth annual Big Data Camp introduced high school juniors and seniors to the advanced tools used by professional data scientists. Campers were required to have coding experience and learned techniques for probing and analyzing massive scientific datasets, such as the dataset from the Array of Things (AoT) urban sensor project. This camp was organized by Argonne's Educational Programs and Outreach staff and taught by ALCF scientists and visualization experts. (Dates: July 25–29, 2022)

Hour of Code

During Computer Science Education Week (CSEdWeek) in December each year, Argonne computer scientists visit Chicagoland schools and assist teachers in celebrating the Hour of Code, a global movement to introduce students to computer programming in a fun way. In 2022, 37 Argonne staff volunteers, 10 from ALCF, visited 44 elementary, middle, and high schools in and around Chicago, to give short tutorials and interact with students on related coding activities, reaching an estimated 2,500+ students. (Dates: December 5–10, 2022)

1.3.3.4 General Outreach

Diversity, Equity, and Inclusion

Through participation in annual Argonne-sponsored outreach events, such as Introduce a Girl to Engineering Day and Science Careers in Search of Women, ALCF staff members connect with young women and introduce them to potential career paths in STEM. ALCF also promotes STEM careers to women through participating in Argonne's Women in Science and Technology program, AnitaB.org's Top Companies for Women Technologists program, and the Grace Hopper Celebration of Women in Computing. ALCF staff also attend the Richard Tapia Celebration of Diversity in Computing Conference to recruit from a diverse set of backgrounds and ethnicities and volunteer as mentors for participants of the National Association for the Advancement of Colored People's (NAACP's) Afro-Academic, Cultural, Technological, and Scientific Olympics (ACT-SO), a year-long achievement program designed to recruit, stimulate, and encourage high academic and cultural achievement among African American high school students.

Facility Tours

Visitors to Argonne can request a tour of the ALCF. ALCF visitors include student groups, Congressional representatives and other government officials, industry representatives, summer research students, visiting researchers, and journalists. Tours are guided by various staff members and can include the machine room, the visualization lab, and the Aurora installation space, where guests have an opportunity to see Aurora being built.

ALCF resumed regular in-person tours in March 2022 while continuing to provide virtual tours for virtual activities, such as training workshops and student summer camps, and when state and county Covid-19 case positivity trends were deemed "high risk" for community transmission. Various staff members hosted more than 80 groups of visitors between March and December 2022. Members of ALCF's leadership team also welcomed U.S. Deputy Secretary of Energy David Turk on August 19, 2022; DOE Office of Science Director Asmeret Asefaw Berhe, who visited Argonne to give the keynote address at the 2022 Postdoctoral Research and Career Symposium on October 25, 2022; and U.S. Secretary of Energy Jennifer Granholm, who came for a ribbon cutting for the Advanced Photon Source (APS) Upgrade on November 4, 2022.

1.3.4 Communications

1.3.4.1 Website Support Center Continuous Improvement

ALCF's online Support Center contains a wide range of resources, from onboarding guides to community announcements to video training tutorials. The Support Center is maintained by

ALCF's media team and technical staff and undergoes internal content reviews, resulting in continuous improvements in the form of new web features and redesigned web pages.

In 2022, the team migrated the ALCF user guides into GitHub. This move aligned with NERSC and OLCF and has made it easier for technical staff to contribute to documentation. It has also moved ALCF from a controlled workflow to a more democratic workflow. Technical staff contributions have increased since this change, which helps to keep information current.

ALCF uses Google Analytics to collect data on how users interact with ALCF's website and has tracked a steady growth in page views across user resources and documentation. ALCF's training webpage attracted 29,187 unique visitors in 2022, an increase of 7,786 from the previous year.

1.3.4.2 Expanding User Experience Across the ALCF

In 2022, ALCF broadened the scope of the documentation committee to include all user-facing touchpoints. The User Experience committee meets every three weeks to work on projects that improve the usability of products and services. The committee provides direct input into user guides, onboarding materials, portal design, application interfaces, and training activities. They also provide feedback on the plans to roll out new machine documentation and services.

With the new production resource Polaris going online in August 2022, ALCF created step-by-step, on-boarding guides for the PIs of all major project awards. Links to these online guides were included in ramp-up emails sent to PIs. This documentation pinpointed user interaction with new components on the Polaris system, which included the transition to using Portable Batch System (PBS) Pro.

1.3.4.3 Consistent Cadence of ALCF Impact on Exascale and AI Efforts

ALCF's communications team continued the two Aurora article series: *Best Practices for GPU Code Development*, highlighting ESP and ECP code optimization efforts for Aurora; and the *Aurora Software Development* series, highlighting the activities and collaborations that are guiding the facility and its users into the next era of scientific computing.

In 2022, ALCF collaborated with Intel on several Aurora Early Science Program articles, "Code Together" podcasts, and Intel product launches; and with SambaNova and Cerebras to jointly promote AI research results on machines in ALCF's AI Testbed, resulting in increased media exposure for ALCF.

1.3.4.4 Communicating Scientific Impact

ALCF produced science stories and articles and promotes HPC training opportunities throughout 2022. Furthermore, ALCF planned marketing campaigns around major annual HPC conferences and events such as the ECP Annual Meeting, International Supercomputing Conference (ISC), Exascale Day, and Supercomputing Conference (SC).

In 2022, ALCF placed 76 original science stories in various news outlets in coordination with Argonne's Communications & Public Affairs (CPA) Division and other ALCF direct relationships. ALCF tracked media hits through its media monitoring service, Meltwater. In 2022, Meltwater reported 1,398 unique ALCF media hits (an increase of 448), 172 of which were chronicled on the ALCF website, and an audience reach of 863.87 million (an increase of 53.21 million). Note: Meltwater defines "reach" as estimating the potential viewership of any article based on the number of visitors to the specific source on both desktop and mobile devices.

ALCF also produced various publications that describe aspects of the facility's mission and summarize its research achievements (Table 1.6). Most of these documents are available for download on the ALCF website.

Table 1.6 Publications Designed for Print

Publication	Frequency	When
Press and Visitor Packets	As Needed	As Needed
Industry Brochure	As Needed	As Needed
Computing Resources	As Needed	As Needed
Annual Report	Yearly	March
Science Report	Yearly	October
Fact Sheets	Yearly	November
INCITE Posters	Yearly	December

1.3.4.5 Messaging for Users and Community

ALCF maintained several communication channels, including direct e-mail campaigns, scriptable e-mail messages, social media postings (Facebook, Twitter, and LinkedIn), and website postings (Table 1.7; target audiences are identified in Table 1.8). Users were able to opt out of the system notify and newsletter mailing lists.

ALCF's monthly e-newsletter, *Newsbytes*, highlighted ALCF-supported research or advancements, promoted training events and allocation program announcements, and linked to relevant news stories. Special announcements about certain training opportunities and fellowships were sent throughout the year, as needed.

Table 1.7 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 1.8 Target Audiences

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

Conclusion

In 2022, ALCF remained focused on ensuring the success of all users, and once again the support metrics and user satisfaction ratings remained at a level expected of a leadership facility. ALCF continued to support research and development teams in adapting their codes to new architectures and helped prepare others to apply for major allocations. ALCF continued to partner with other national laboratories and the ECP and to present work in premier scientific journals and at professional meetings and conferences. The Introduction to AI-driven Science on Supercomputers training series was held for a second year with great success, attracting more than 200 attendees from 90 universities. During 2022, ALCF users also gained access to the new Polaris system and the AI Testbed (SambaNova and Cerebras). Lastly, improvements to ALCF's project management software, Userbase3, have further streamlined the user account process.

Section 2. Operational Performance

Did the facility's operational performance meet established targets?

2.1 ALCF Response

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

Tables 2.1 and 2.2 summarize all operational performance metrics of HPC computational and storage systems reported in this section.

Table 2.1 Summary of Operational Performance of HPC Computational Systems

	Theta (with expansion) Theta (Cray XC40):4008-node, 251K-core 64 TB MCDRAM 770 TB DDR4 Theta expansion (NVIDIA DGX): 24-node 24 TB of DDR4 RAM 7.68 TB of GPU memory				Polaris (HPE Apollo 6500 Gen 10+) 560-node, 17920-core 2240 NVIDIA A100 GPU 287 TB of DDR4 RAM 90 TB HBM2	
	CY 2021		CY 2022		CY 2022	
	Target	**Actual**	**Target**	**Actual**	**Target**	**Actual**
Scheduled Availability	90.0%	99.40%	90.00%	98.6%	90.00%	98.2%
Overall Availability	90.0%	95.10%	90.00%	95.7%	90.00%	94.7%
System MTTI	N/A [e]	12.07 days	N/A	10.62 days	N/A	9.86 days
System MTTF	N/A	45.51 days	N/A	32.87 days	N/A	47.78 days
Expansion System MTTI	N/A	13.46 days	N/A	12.53 days	N/A	N/A
Expansion System MTTF	N/A	121.54 days	N/A	60.39 days	N/A	N/A
INCITE Usage	17.8M	20.8M	17.8M	20.4M	1322.0K	589.6K
Total Usage	N/A	33.5M	N/A	33.2M	N/A	1163.6K
System Utilization	N/A	98.1%	N/A	97.0%	N/A	67.3%
INCITE Overall Capability [a,b]	20.0%	84.8%	20.0%	85.8%	20.0%	47.7%
INCITE High Capability [c,d]	N/A	22.7%	N/A	27.2%	N/A	22.3%

[a] Polaris Overall Capability = Jobs using ≥ 20.0 percent (99 nodes).

[b] Theta (with expansion) Overall Capability = Jobs using ≥ 20.0 percent (800 nodes) of Theta.

[c] Polaris High Capability = Jobs using ≥ 60.0 percent (297 nodes).

[d] Theta (with expansion) High Capability = Jobs using ≥ 60.0 percent (2,400 nodes) of Theta.

[e] N/A = not applicable.

Table 2.2 Summary of Operational Performance of HPC Storage Systems

Theta-fs0 File System Cray Sonexion 2000 with 9.2 PB of usable storage				
	CY 2021		CY 2022	
	Target	Actual	Target	Actual
Scheduled Availability	90.0%	100.0%	90.0%	99.3%
Overall Availability	90.0%	96.1%	90.0%	96.3%
System MTTI	N/A [a]	15.25 days	N/A	13.52 days
System MTTF	N/A	182.47 days	N/A	72.51 days
Grand File System Cray E1000 with 100 PB of storage at 650 GB/s				
	CY 2021		CY 2022	
	Target	Actual	Target	Actual
Scheduled Availability	N/A	100.0%	90.0%	99.5%
Overall Availability	N/A	96.5%	90.0%	96.6%
System MTTI	N/A	15.31 days	N/A	14.68 days
System MTTF	N/A	182.45 days	N/A	121.13 days
Eagle File System Cray E1000 with 100 PB of storage at 650 GB/s				
	CY 2021		CY 2022	
	Target	Actual	Target	Actual
Scheduled Availability	N/A	100.0%	90.0%	99.5%
Overall Availability	N/A	96.5%	90.0%	96.6%
System MTTI	N/A	16.01 days	N/A	14.68 days
System MTTF	N/A	365.00 days	N/A	121.13 days
HPSS Archive LTO8 tape drives and tape with 350 PB of storage capacity				
	CY 2021		CY 2022	
	Target	Actual	Target	Actual
Scheduled Availability	90.0%	100.0%	90.0%	99.7%
Overall Availability	90.0%	96.5%	90.0%	96.7%
System MTTI	N/A	16.01 days	N/A	14.70 days
System MTTF	N/A	365.00 days	N/A	121.28 days

[a] N/A = not applicable.

2.2 ALCF Production Resources Overview

During CY 2022, the ALCF operated several production resources.

- Polaris is a 560 node, ~18K core, HPE Apollo 6500 Gen 10+ with 287 TB of RAM and 2240 NVIDIA A100 GPUs. Polaris went into production on August 9, 2022.
- Theta is a 4,392-node, ~281K-core, 11.69-PF Cray XC40 with 892 TB of RAM.
- ThetaGPU is a 24-node expansion to Theta consisting of an NVIDIA DGX A100-based system. The DGX A100 comprises eight NVIDIA A100 GPUs with AMD EPYC 7742 CPUs for a total of 24 TB of DDR4 RAM and 7.68 TB of GPU memory.
- Grand and Eagle are each 100PB Lustre file systems and are mounted facility wide. The only difference is that Eagle has Globus sharing enabled and is used as a “community file system.”
- The facility wide high-performance storage system (HPSS) tape archive is comprised of three 10,000-slot libraries with LTO8 drives and tapes, with some legacy LTO6 drives and tapes. Currently, the tape libraries have a maximum storage capacity of 305 PB.

2.3 Definitions

- *Overall availability* is the percentage of time a system is available to users. Outage time reflects both scheduled and unscheduled outages.
- *Scheduled availability*, for HPC Facilities, 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 (for 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 at least seven calendar days prior. If a 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 if the return to service is four or more hours later than the scheduled end time. For storage resources, availability will be if any user can read and write any portion of the disk space. The availability metric provides measures that are indicative of the stability of the systems and the quality of the maintenance procedures.
- *Mean time to interrupt (MTTI)* is the time, on average, to any outage on the system, whether unscheduled or scheduled. Also known as MTBI (Mean Time Between Interrupt).
- *Mean time to failure (MTTF)*, is the time, on average, to an unscheduled outage on the system.
- *Usage* is defined as resources consumed in units of node-hours.

- *Utilization* is the percentage of the available node-hours used (i.e., a measure of how busy the system was kept when it was available).
- *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 category, including staff and vendors.

For more information on performance metric calculations, see Appendix A.7.

2.4 Polaris

2.4.1 Scheduled and Overall Availability

Polaris entered full production on August 9, 2022. In consultation with ALCF's DOE Program Manager, ALCF has agreed to a target of 90 percent overall availability and a target of 90 percent scheduled availability. (ASCR requested that all user facilities use a target of 90 percent for scheduled availability for the lifetime of the production resources). Table 2.3 summarizes the Polaris availability.

Table 2.3 Availability Results

Polaris (HPE Apollo 6500 Gen 10+) 560-node, 17920-core 2240 NVIDIA A100 GPU 287 TB DDR4 0 TB HBM2		
	CY 2022	
	Target (%)	Actual (%)
Scheduled Availability	90.0	98.2
Overall Availability	90.0	94.7

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.

2.4.1.1 Explanation of Significant Availability Losses

This section briefly describes the causes of major losses of availability for the period August 9, 2022, through December 31, 2022, as shown in Figure 2.1.

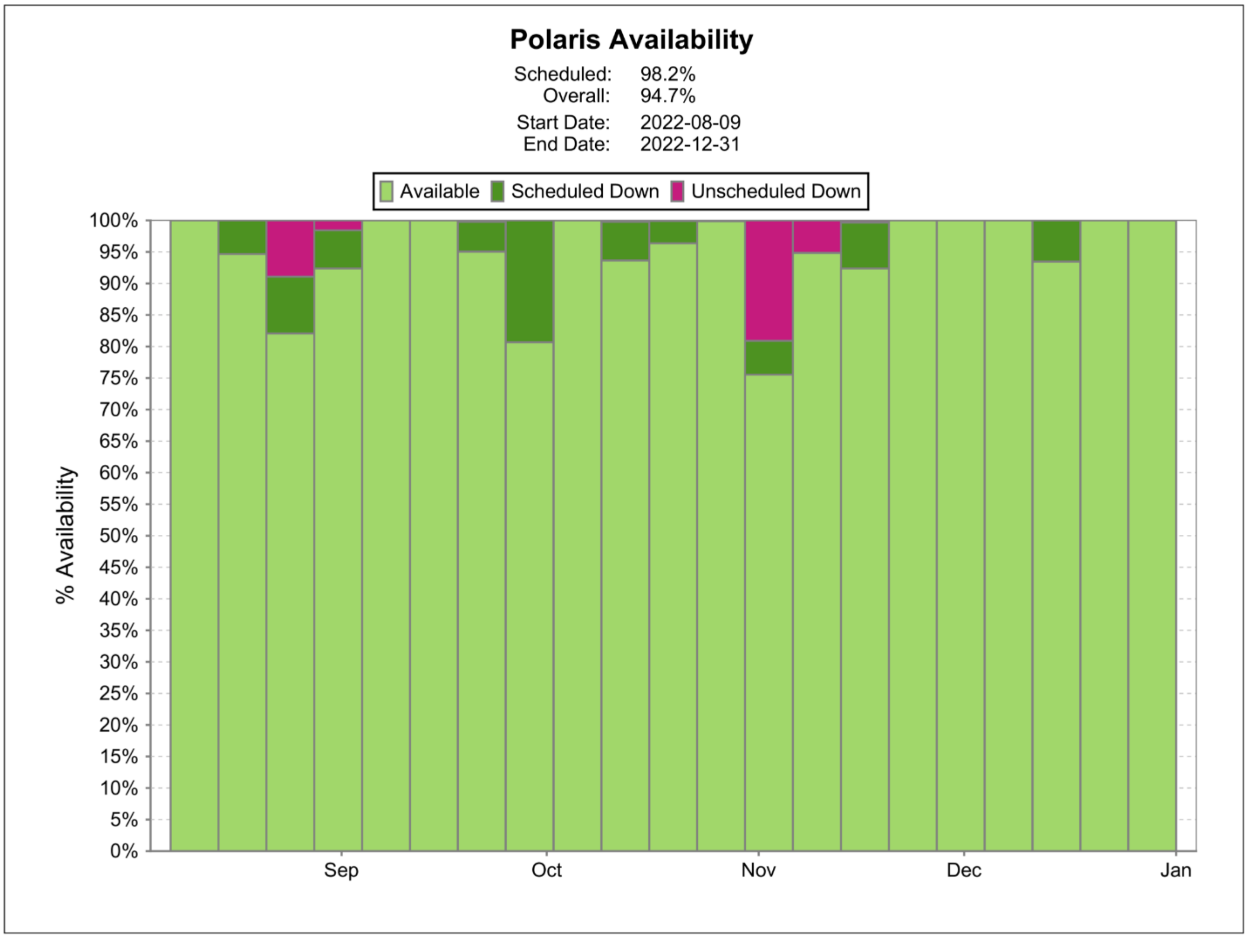

Figure 2.1 Polaris Weekly Availability for CY 2022

Graph Description: Each bar in Figure 2.1 represents the percentage of the machine available for seven days. Each bar accounts for all of the time in one of three categories. The pale-green portion represents available node-hours; the darker green represents scheduled downtime for that week; and magenta represents unscheduled downtime. Significant loss events are described in detail below.

August 25, 2022: Partial Unscheduled outage – pbs license
Coming out of scheduled maintenance, 499 nodes were missing their PBS license. PBS did not schedule jobs on those nodes until the license issue was resolved.

October 31, 2022: Unscheduled outage – image issues
ALCF extended the scheduled outage due to reboot issues. Polaris nodes were running out of file handles. Resolution: increase the file handle limit to 500K.

November 8, 2022: Unscheduled outage scheduler issue
This was due to a bug that occurred if a user entered the same job ID twice on a qdel command line, it would cause the scheduler to segfault. The scheduler would restart, but the server would hold the handle to the previous instance, so nothing would schedule.

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

2.4.2.1 MTTI and MTTF Summary

MTTI and MTTF are reportable values with no specific targets. Table 2.4 summarizes the current MTTI and MTTF values, respectively, for Polaris.

Table 2.4 MTTI and MTTF Results

Polaris (HPE Apollo 6500 Gen 10+) 560-node, 17920-core 2240 NVIDIA A100 GPU 287 TB DDR4 90 TB HBM2		
	CY 2022	
	Target	**Actual**
System MTTI	N/A [a]	9.86 days
System MTTF	N/A	47.78 days

[a] N/A = not applicable.

Polaris currently has a biweekly maintenance schedule to perform upgrades, hardware replacements, OS upgrades, etc. ALCF uses these preventative maintenance (PM) opportunities to schedule other potentially disruptive maintenance such as facility power and cooling work and storage system upgrades and patching. Although Polaris's biweekly maintenance schedule does not directly affect MTTF, it generally tends to cap MTTI at 14 days.

2.4.3 Resource Utilization

The following sections discuss system allocation and usage, system utilization percentage, and capability usage.

2.4.3.1 System Utilization

System 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. Polaris utilization for its initial 5 months was below 80 percent; however, figure 2.2 shows utilization increasing above 80 percent after the first 2 months. Table 2.5 summarizes ALCF system utilization results, and Figure 2.2 shows system utilization over time by program.

Table 2.5 System Utilization Results

Polaris (HPE Apollo 6500 Gen 10+) 560-node, 17920-core 2240 NVIDIA A100 GPU 287 TB DDR4 90 TB HBM2		
	CY 2022	
	Target	**Actual**
System Utilization	N/A [a]	67.3%

[a] N/A = not applicable.

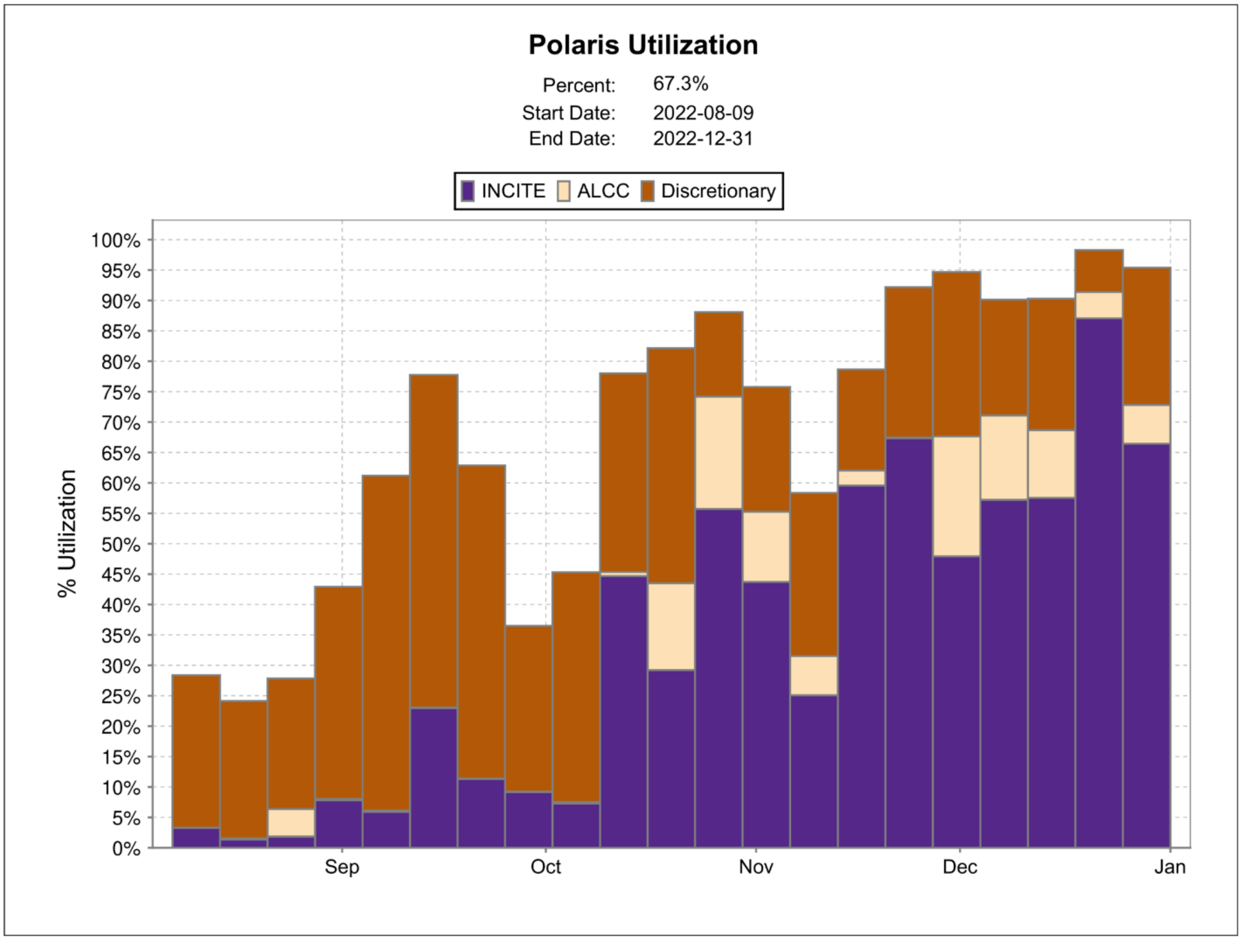

Figure 2.2 Polaris System Utilization over Time by Program

The system utilization for Polaris was 67.3 percent for its 2022 production period of August 9, 2022, through December 31, 2022.

Table 2.6 shows how Polaris's system hours were allocated and used by allocation source. Most DD projects are exploratory investigations, so the allocations are often not used in full.
DD allocations are discussed in detail in Section 3.1.2. In CY 2022, Polaris delivered a total of 1,163,601 node-hours.

Table 2.6 Node-Hours Allocated and Used by Program

Polaris (HPE Apollo 6500 Gen 10+) 560-node 17920-core 2240 Nvidia A100 GPU 287 TB DDR4 90 TB HBM2			
	CY 2022		
	Allocated	Used	
	Node-hours	Node-hours	%
INCITE	1322.0K	589.6K	50.7%
ALCC	355.8K	92.6K	8.0%
DD [a]	543.6K	481.5K	41.3%
Total	2221.4K	1163.6K	100.0%

[a] DD node-hours include ESP node-hours.

Summary: For CY 2022, the system usage and system utilization values were in line with general expectations. The calculations for system utilization are described in Appendix A.

2.4.3.2 Capability Utilization

Polaris has a total of 560 nodes of which 54 were purchased as extra nodes and 8 are debug nodes. Therefore, ALCF uses 496 nodes to calculate capability metrics. On Polaris, capability is defined as using greater than 20 percent of the 496 nodes, or 99 nodes, and high capability is defined as using greater than 60 percent of the 496 nodes, or 297 nodes. See Table A.2 in Appendix A for more detail on the capability calculation. Table 2.7 and Figure 2.3 show that ALCF has substantially exceeded 20 percent capability on Polaris for INCITE and all other programs. Figure 2.4 shows the three programs' utilization of total core hours (from Table 2.7) over time, and Figure 2.5 shows the overall distribution of job sizes over time.

Table 2.7 Capability Results

Polaris (HPE Apollo 6500 Gen 10+) 560-node, 17920-core 2240 NVIDIA A100 GPU 287 TB DDR4 90 TB HBM2			
	CY 2022		
Capability Usage	**Total Hours**	**Capability Hours**	**Percent Capability**
INCITE Overall [a]	589.5K	281.1K	47.7%
INCITE High [b]	589.6K	131.5K	22.3%
ALCC Overall	92.6K	34.7K	37.5%
ALCC High	92.6K	3.6K	3.9%
Director's Discretionary Overall [c]	383.7K	160.9K	41.9%
Director's Discretionary High [c]	383.7K	15.4K	4.0%
ESP Director's Discretionary Overall	97.8K	39.4K	40.3%
ESP Director's Discretionary High	97.8K	1.9K	2.0%
TOTAL Overall	1163.6K	516.1K	44.4%
TOTAL High	1163.6K	152.4K	13.1%

[a] Polaris Overall Capability = Jobs using ≥ 20.0 percent (99 nodes).

[b] Polaris High Capability = Jobs using ≥ 60.0 percent (297 nodes).

[c] Does not include ESP.

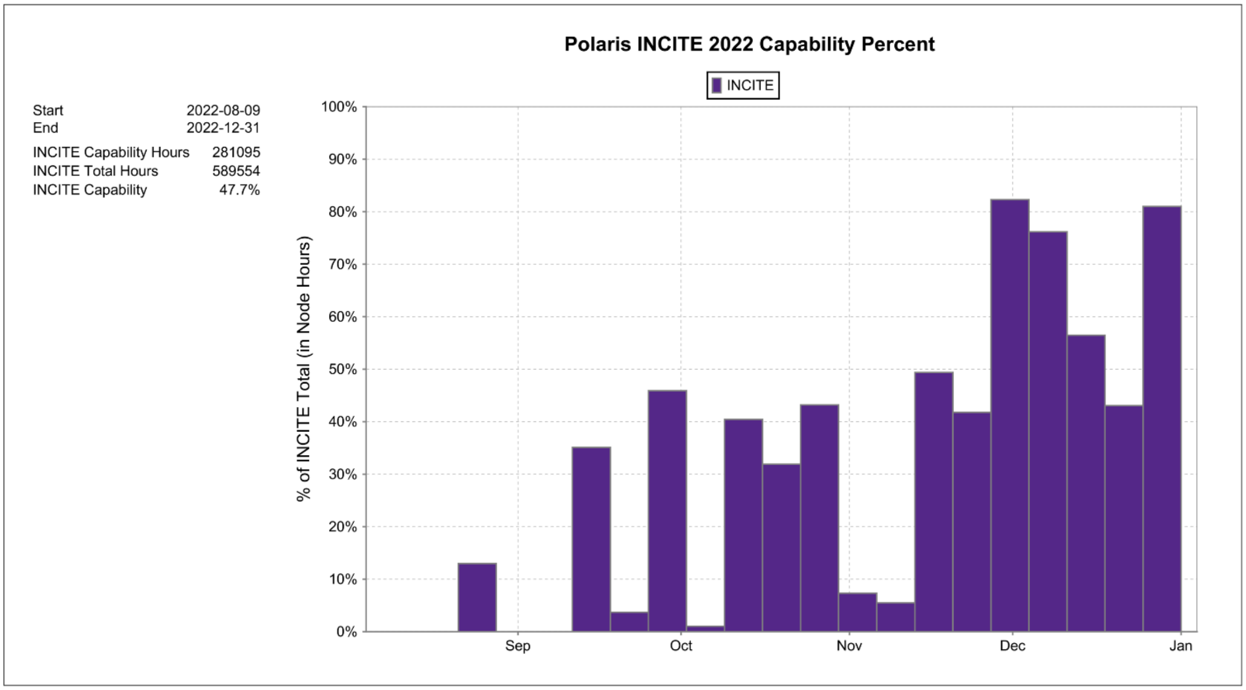

Figure 2.3 Polaris INCITE Overall Capability

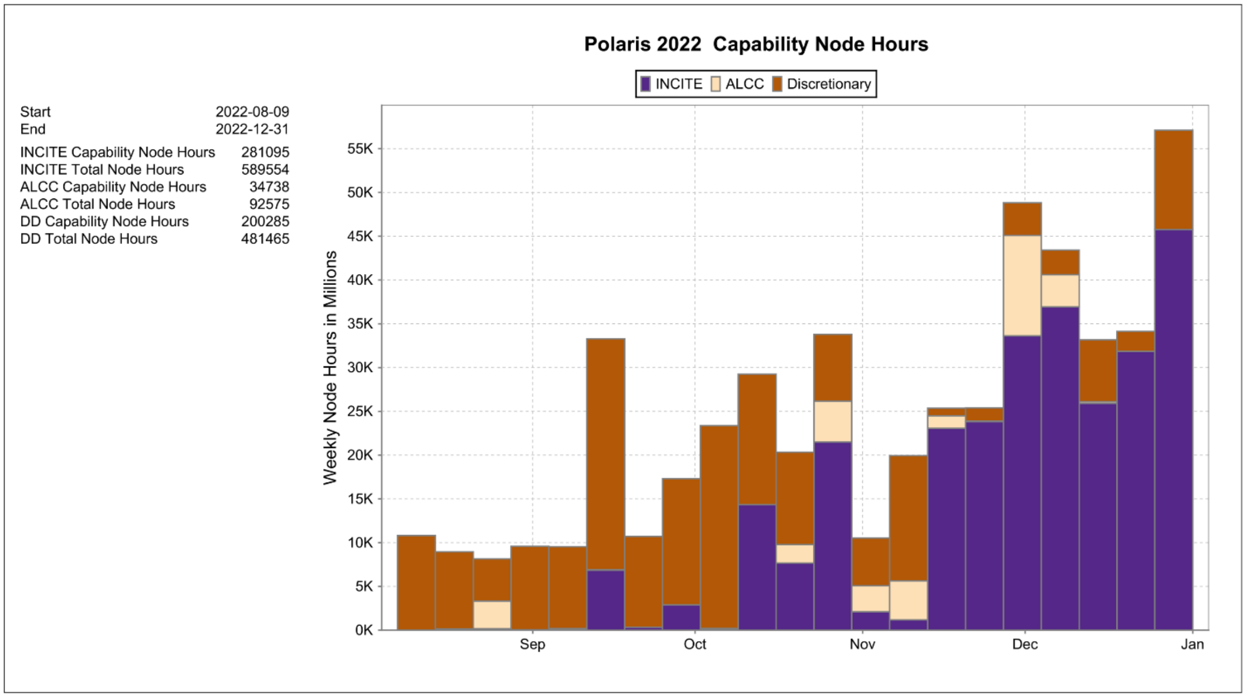

Figure 2.4 Polaris Capability Node-Hours by Program

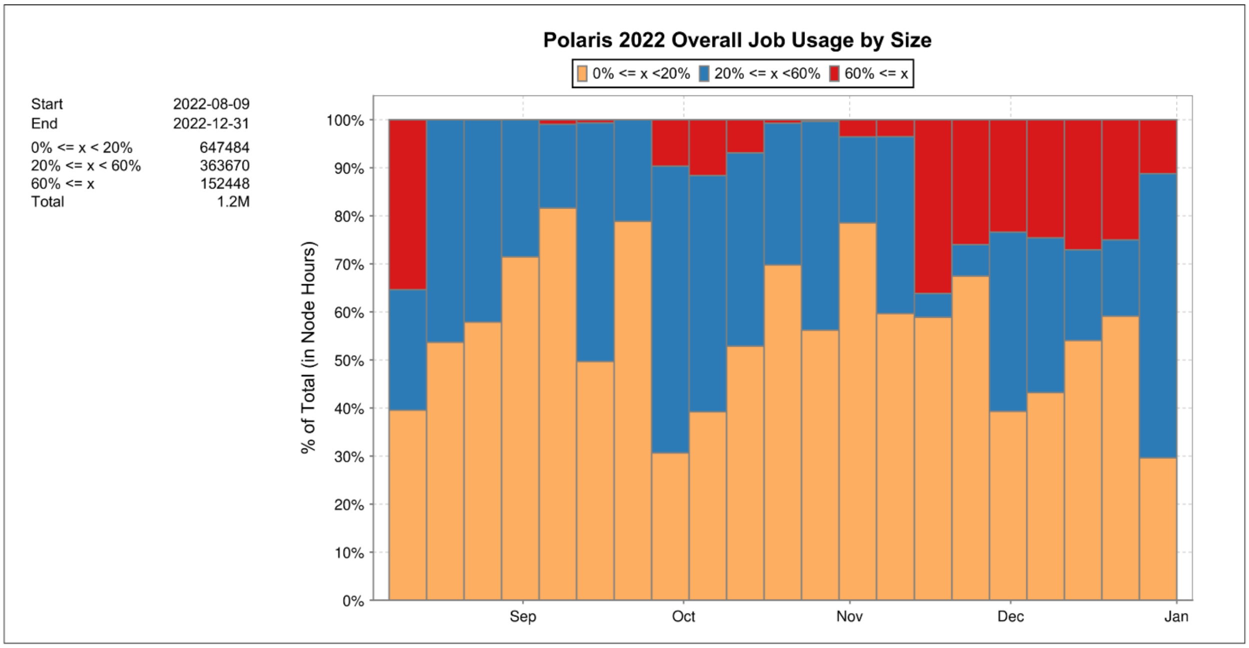

Figure 2.5 Polaris Job Usage by Size

2.5 Theta

2.5.1 Scheduled and Overall Availability

Theta entered full production on July 1, 2017. The GPU expansion to Theta entered production on January 1, 2021. In consultation with ALCF's DOE Program Manager, ALCF has agreed to a target of 90 percent overall availability and a target of 90 percent scheduled availability (ASCR requested that all user facilities use a target of 90 percent for scheduled availability for the lifetime of the production resources). Table 2.8 summarizes the availability results for Theta (with expansion).

Table 2.8 Availability Results

Theta (with expansion) Theta (Cray XC40):4008-node, 251K-core 64 TB MCDRAM 770 TB DDR4 Theta expansion (NVIDIA DGX): 24-node 24 TB of DDR4 RAM 7.68 TB of GPU memory				
	CY 2021		CY 2022	
	Target (%)	Actual (%)	Target (%)	Actual (%)
Scheduled Availability	90.0	99.4	90.0	98.6
Overall Availability	90.0	95.1	90.0	95.7

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.

2.5.1.1 Explanation of Significant Availability Losses

This section briefly describes the causes of major losses of availability for the period January 1, 2022, through December 31, 2022, as shown in Figure 2.6.

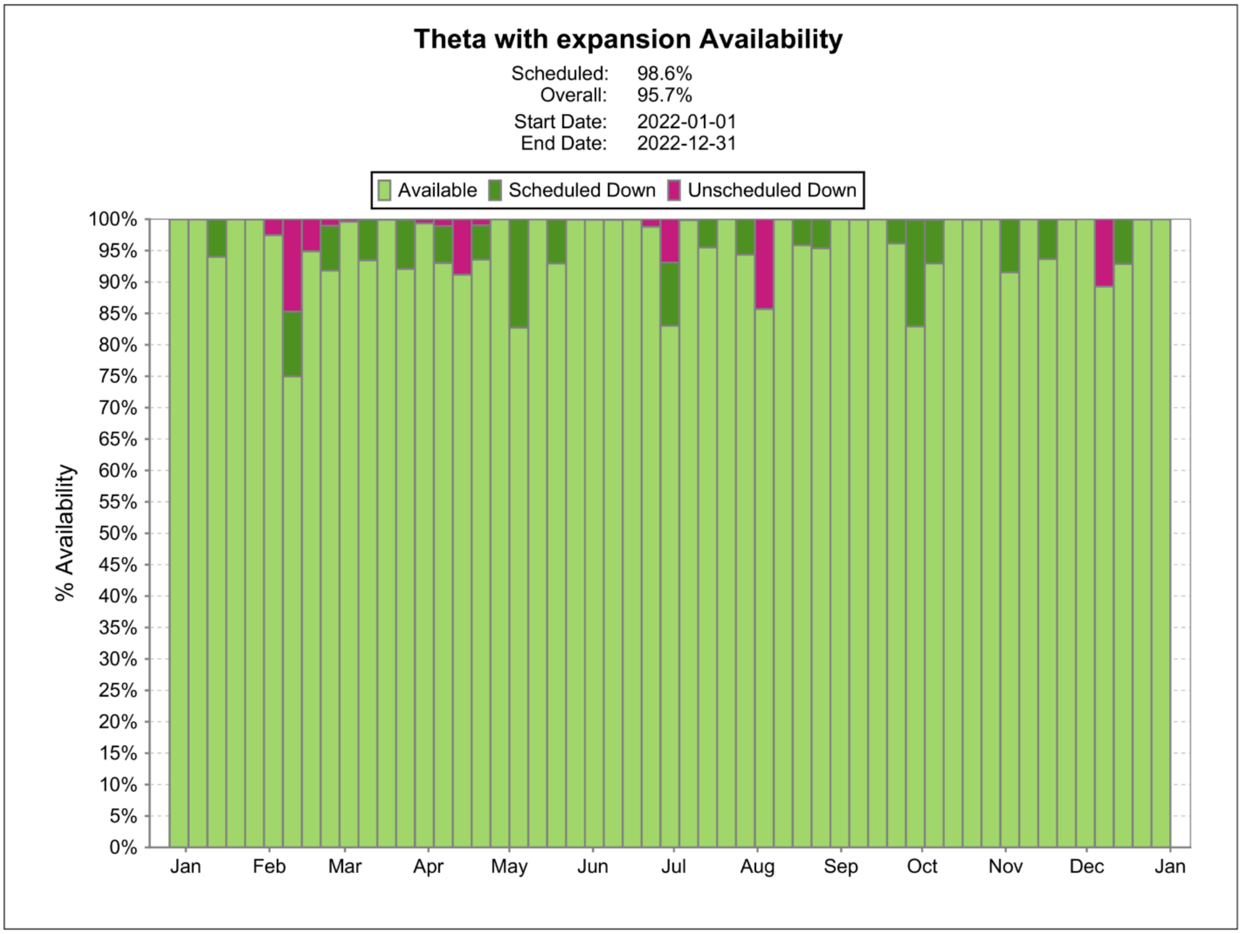

Figure 2.6 Theta (with expansion) Weekly Availability for CY 2022

Graph Description: Each bar in Figure 2.6 represents the percentage of the machine available for seven days. Each bar accounts for all the time in one of three categories. The pale-green portion represents available node-hours; the darker green represents scheduled downtime for that week; and magenta represents unscheduled downtime. Significant loss events are summarized below.

February 4, 2022: Unscheduled outage – LNET errors
File access errors affected jobs accessing files on theta-fs0.

February 12, 2022: Unscheduled outage – Theta-fs0 file system outage
The file system, theta-fs0 was inaccessible. Scheduling was halted during the file system recovery.

February 14, 2022: Unscheduled outage – Eagle file system
The Eagle file system was partially inaccessible. Scheduling was halted during the file system recovery.

April 12, 2022: Unscheduled outage – Scheduling paused
The scheduler on theta was paused due to Eagle file system issues.

April 13, 2022: Unscheduled outage – Scheduler failure
The Cobalt scheduler received corrupted data from the ALPS system. The scheduler was restarted, and an issue was reported to HPE.

April 16, 2022: Unscheduled outage – Grand file system issue
Two drives on the Grand file system failed and the scheduler was paused. After the Grand issue was resolved, issues arose on the Aries network, which were resolved by rebooting the affected nodes.

June 20, 2022: Unscheduled outage – jobs stuck in kill state
Multiple jobs were stuck in the kill state due to two file system nodes that failed. The outage was only 2 hours. Resolution: reboot of the nodes after the file system issue was resolved.

June 28, 2022: Unscheduled outage – home file system migration
Due to the delay in the return to service of Theta and ThetaGPU after a scheduled maintenance, an unscheduled outage occurred. Theta and ThetaGPU were withheld until issues related to the home filesystem migration were resolved.

August 1, 2022: Unscheduled outage – cooling system
All ALCF systems were shut down due to a cooling system failure. A faulty hose connection resulted in one of the chillers shutting down. The remaining chiller did not have enough capacity and temperatures in the data center rose rapidly. ALCF quickly shut down all hardware to avoid a catastrophic hardware loss. Within two days, the cooling system was repaired and the systems were restarted.

December 7, 2022: Unscheduled outage – Theta-fs0 outage
The file system theta-fs0 failed during a hot swap of failed components. The hot swaps failed due to mismatched certs being pushed from the configuration management system.

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

2.5.2.1 MTTI and MTTF Summary

MTTI and MTTF are reportable values with no specific targets. Table 2.9 summarizes the current MTTI and MTTF values, respectively, for Theta. Because Theta with expansion is an NVIDIA DGX system and Theta is a Cray XC40 system, MTTI and MTTF are calculated and reported separately for each.

Table 2.9 MTTI and MTTF Results

Theta (Cray XC40):4008-node, 251K-core 64 TB MCDRAM 770 TB DDR4				
	CY 2021		CY 2022	
	Target	**Actual**	**Target**	**Actual**
System MTTI	N/A [a]	12.07 days	N/A	10.62 days
System MTTF	N/A	45.51 days	N/A	32.87 days

Theta expansion (NVIDIA DGX): 24-node 24 TB of DDR4 RAM 7.68 TB of GPU memory				
	CY 2021		CY 2022	
	Target	**Actual**	**Target**	**Actual**
System MTTI	N/A	13.46 days	N/A	12.53 days
System MTTF	N/A	121.54 days	N/A	60.39 days

[a] N/A = not applicable.

Theta currently functions on a biweekly maintenance schedule to perform Cray driver upgrades, hardware replacements, OS upgrades, etc. ALCF uses these preventative maintenance (PM) opportunities to schedule other potentially disruptive maintenance such as facility power and cooling work and storage system upgrades and patching. Although Theta's biweekly maintenance schedule does not directly affect MTTF, it generally tends to cap MTTI at 14 days.

2.5.3 Resource Utilization

The following sections discuss system allocation and usage, system utilization percentage, and capability usage.

2.5.3.1 System Utilization

System 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 2.10 summarizes ALCF system utilization results, and Figure 2.7 shows system utilization over time by program.

Table 2.10 System Utilization Results

Theta (with expansion) Theta (Cray XC40):4008-node, 251K-core 64 TB MCDRAM 770 TB DDR4 Theta expansion (NVIDIA DGX): 24-node 24 TB of DDR4 RAM 7.68 TB of GPU memory				
	CY 2021		CY 2022	
	Target	Actual	Target	Actual
System Utilization	N/A [a]	98.1%	N/A	97.0%

[a] N/A = not applicable.

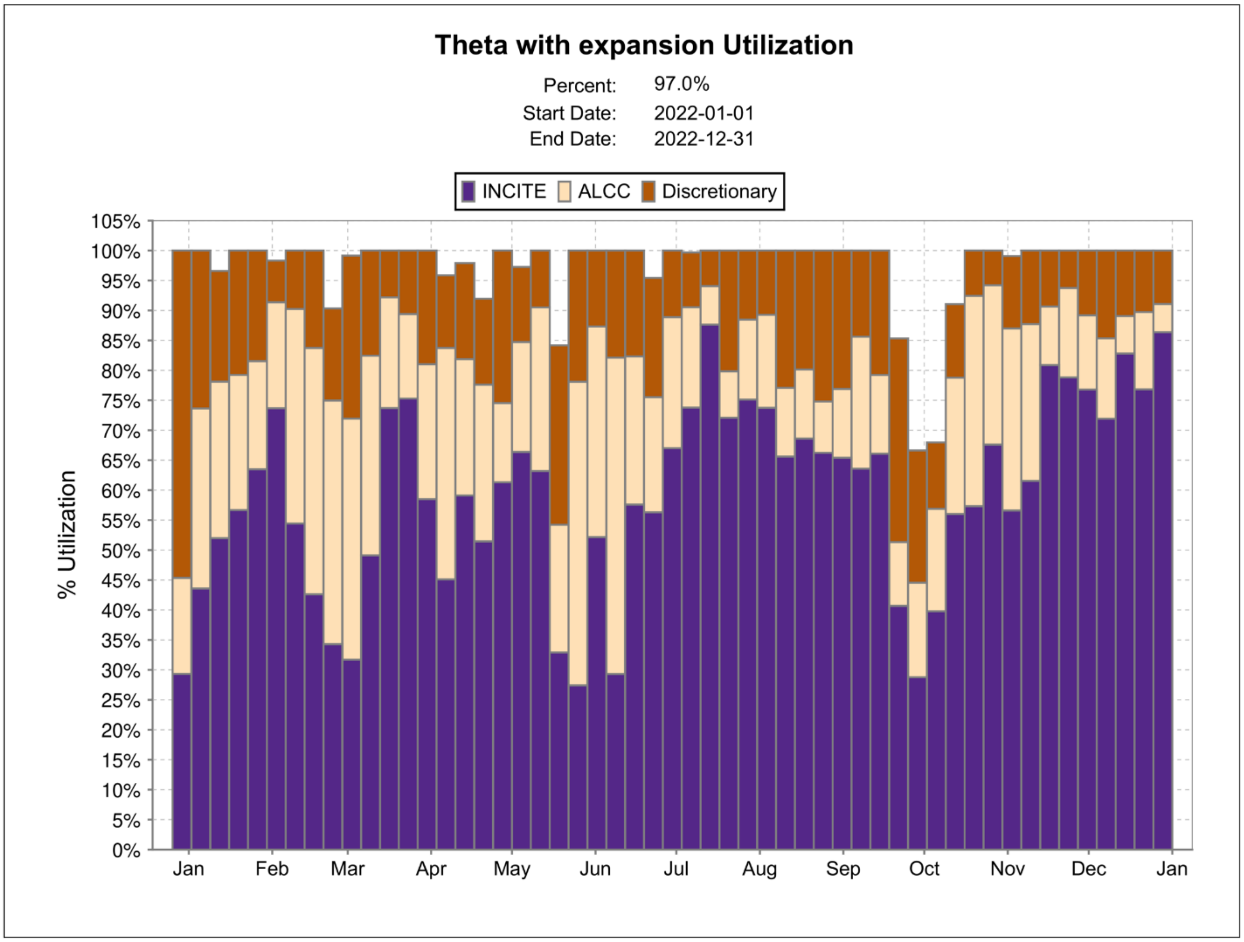

Figure 2.7 System Utilization over Time by Program

The system utilization for Theta was 97.0 percent for its 2022 production period of January 1, 2022, through December 31, 2022.

Table 2.11 shows how Theta's system hours were allocated and used by allocation source. Multiplying the theoretical hours by availability and system utilization values that were agreed upon with ALCF's DOE Program Manager determines the hours available. Of the hours available, 60 percent were allocated to the INCITE program, up to 20 percent to ALCC program allocations, and 20 percent to DD allocations. The allocated values for the DD allocations appear higher than expected because they represent a rolling allocation. Most DD projects are exploratory investigations, so the time allocations are often not used in full. DD allocations are discussed in detail in Section 3.1.2. In CY 2022, Theta (with expansion) delivered a total of 33.2 million node-hours.

Table 2.11 Node-Hours Allocated and Used by Program

Theta (with expansion) Theta (Cray XC40):4008-node, 251K-core 64 TB MCDRAM 770 TB DDR4 Theta expansion (NVIDIA DGX): 24-node 24 TB of DDR4 RAM 7.68 TB of GPU memory						
	CY 2021			CY 2022		
	Allocated	Used		Allocated	Used	
	Node-hours	Node-hours	%	Node-hours	Node-hours	%
INCITE	17.8M	20.8M	62.1	17.8M	20.4M	61.5
ALCC	7.3M	7.2M	21.6	6.4M	7.4M	7.9
DD	7.8M	5.5M	16.3	7.1M	5.3M	16.0
Total	32.9M	33.5M	100.0	31.3M	33.2M	100.0

Summary: For CY 2022, the system usage and system utilization values were in line with general expectations. The calculations for system utilization are described in Appendix A.

2.5.3.2 Capability Utilization

For Theta, capability is defined as using greater than 20 percent of the machine, or 800 nodes, and high capability is defined as using greater than 60 percent of the machine, or 2,400 nodes. See Table A.2 in Appendix A for more detail on the capability calculation. Table 2.12 and Figure 2.8 show that ALCF has substantially exceeded the 20 percent capability metric set for INCITE. The data for the ALCC and DD projects are provided as a reference even though no capability targets were defined. Figure 2.9 shows the three programs' utilization of total core hours (from Table 2.12) over time, and Figure 2.10 shows the overall distribution of job sizes over time.

Table 2.12 Capability Results

Theta (with expansion) Theta (Cray XC40):4008-node, 251K-core 64 TB MCDRAM 770 TB DDR4 Theta expansion (NVIDIA DGX): 24-node 24 TB of DDR4 RAM 7.68 TB of GPU memory						
	CY 2021			CY 2022		
Capability Usage	**Total Hours**	**Capability Hours**	**Percent Capability**	**Total Hours**	**Capability Hours**	**Percent Capability**
INCITE Overall [a]	20.8M	17.6M	84.8%	20.4M	17.5M	85.8%
INCITE High [b]	20.8M	4.7M	22.7%	20.4M	5.6M	27.2%
ALCC Overall	7.2M	3.0M	41.1%	7.4M	4.6M	61.7%
ALCC High	7.2M	0.6M	7.7%	7.4M	0.5M	6.6%
Director's Discretionary Overall	5.5M	1.1M	20.9%	5.3M	0.8M	15.5%
Director's Discretionary High	5.5M	0.2M	4.2%	5.3M	0.1M	1.8%
TOTAL Overall	33.1M	24.1M	72.9%	33.2M	22.9M	69.1%
TOTAL High	33.1M	6.1M	18.4%	33.2M	6.1M	18.5%

[a] Theta (with expansion) Overall Capability = Jobs using ≥ 20.0 percent (800 nodes) of Theta.

[b] Theta (with expansion) High Capability = Jobs using ≥ 60.0 percent (2400 nodes) of Theta.

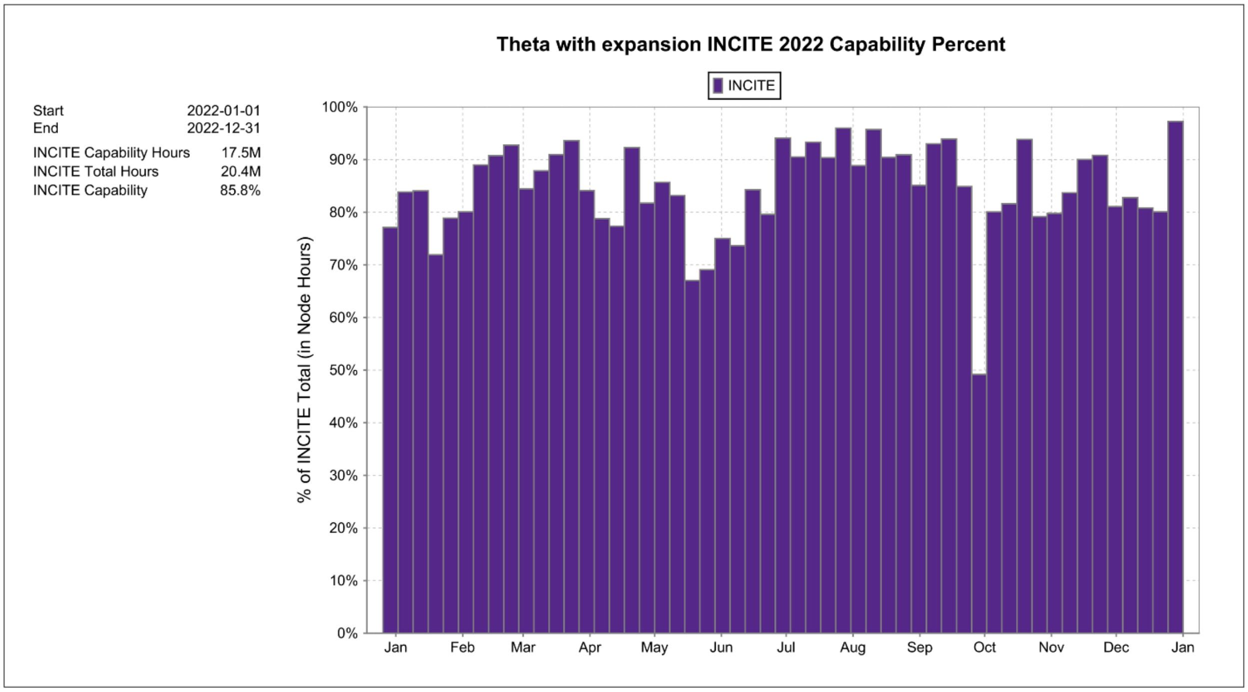

Figure 2.8 Theta INCITE Overall Capability

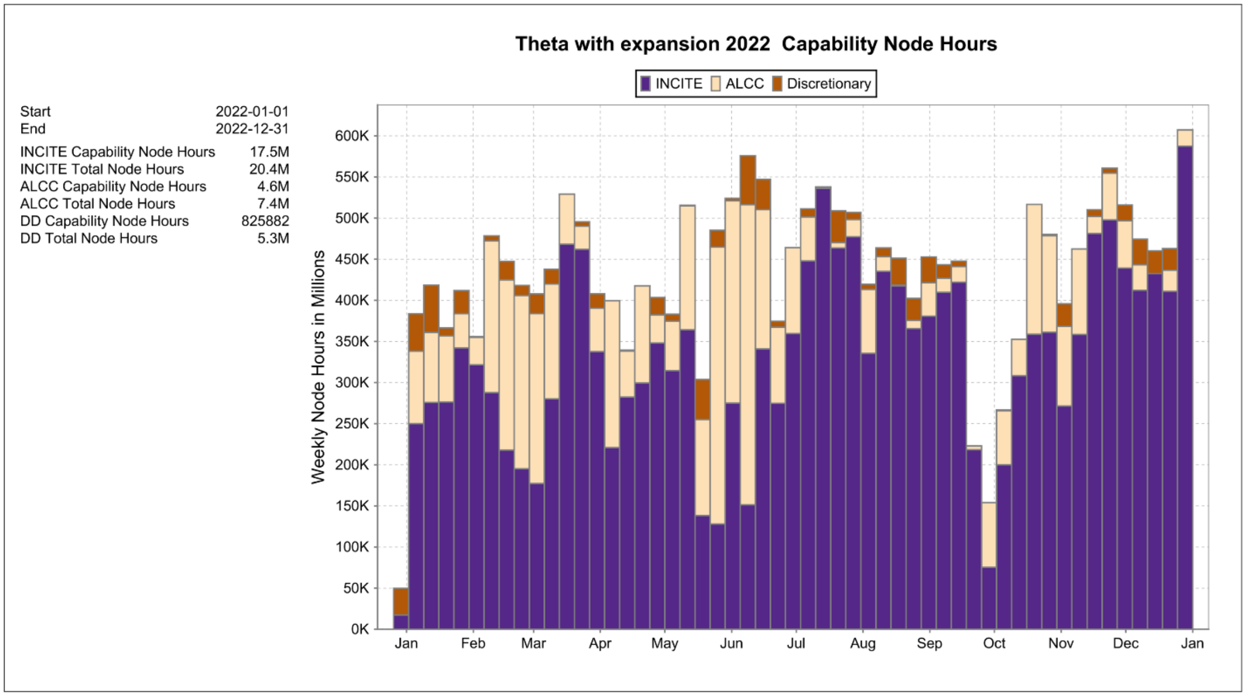

Start	2022-01-01
End	2022-12-31
INCITE Capability Node Hours	17.5M
INCITE Total Node Hours	20.4M
ALCC Capability Node Hours	4.6M
ALCC Total Node Hours	7.4M
DD Capability Node Hours	825882
DD Total Node Hours	5.3M

Figure 2.9   Theta Capability Node-Hours by Program

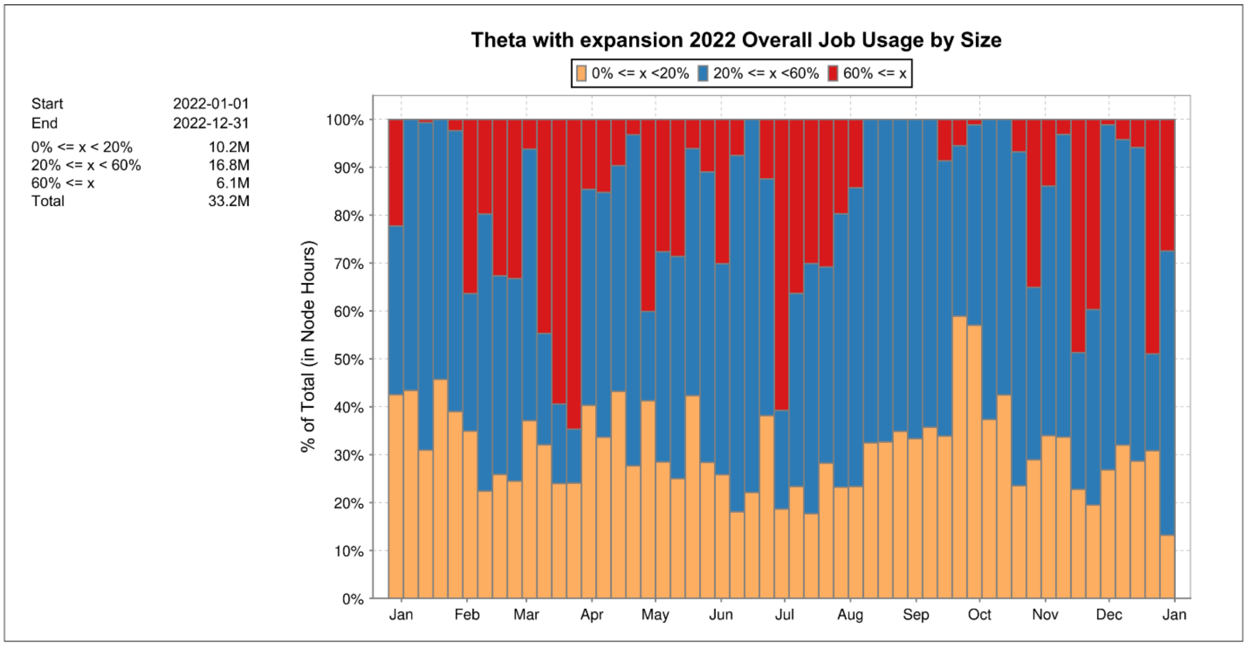

Start	2022-01-01
End	2022-12-31
0% <= x < 20%	10.2M
20% <= x < 60%	16.8M
60% <= x	6.1M
Total	33.2M

Figure 2.10   Theta Job Usage by Size

2.6 Storage

This section covers availability and MTTI/F metrics for the production storage resources.

2.6.1 Theta-fs0 Lustre File System

Theta-fs0 is a Cray Sonexion 3000 Lustre file system with 9.2 PB of usable space that is mounted by Theta, ThetaGPU, and the data and analysis resource. Theta and Theta-fs0 were installed together, and both entered production on July 1, 2017. Theta and Theta-fs0 were not treated as separate entities in previous OARs.

2.6.1.1 Scheduled and Overall Availability

ALCF used the target metrics of 90 percent overall availability scheduled availability as proposed in the CY 2021 OAR. Theta-fs0 is tightly tied to Theta, so this follows ASCR’s request that all user facilities use a target of 90 percent for scheduled availability for the lifetime of a production resource. Table 2.13 summarizes the availability results.

Table 2.13 Availability Results

Theta-fs0 File System Cray Sonexion 3000 with 9.2 PB of usable storage				
	CY 2021		CY 2022	
	Target (%)	Actual (%)	Target (%)	Actual (%)
Scheduled Availability	N/A [a]	100.0	90	99.3
Overall Availability	N/A	96.1	90	96.3

[a] N/A = not applicable.

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 the Theta-fs0 file system for the period of January 1, 2022, through December 31, 2022, as noted in Figure 2.11.

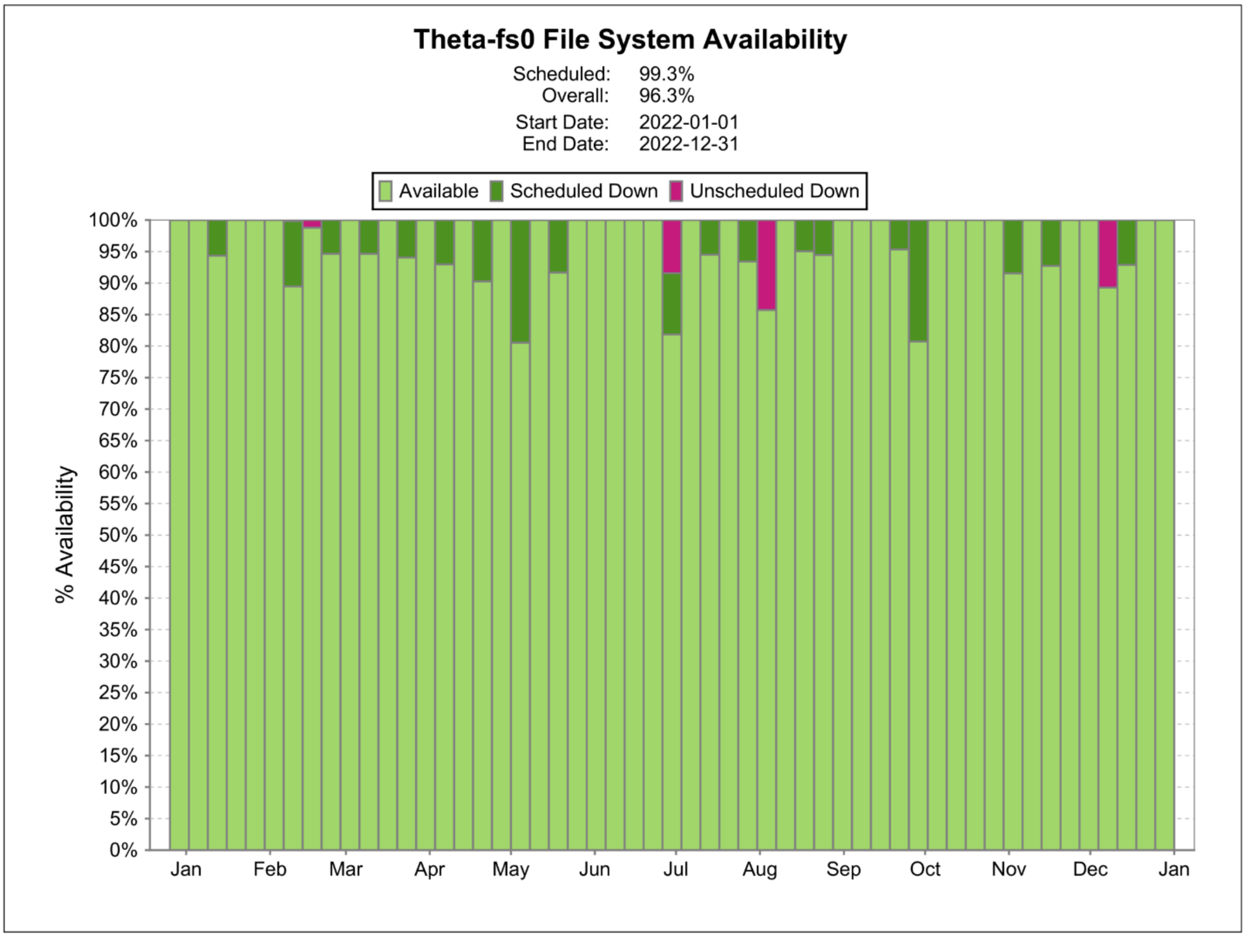

Figure 2.11 Theta-fs0 Weekly Availability for CY 2022

Graph Description: Each bar in Figure 2.11 represents the percentage of the file system available for seven days. Each bar accounts for all the time in one of three categories. The pale-green portion represents available node-hours; the darker green represents scheduled downtime for that week; and magenta represents unscheduled downtime. Each of the significant loss events is described in detail below (these also appeared in the Theta section).

February 4, 2022: Unscheduled outage – LNET errors
File access errors affected jobs accessing files on theta-fs0.

February 12, 2022: Unscheduled outage – Theta-fs0 file system outage
The file system, theta-fs0 was inaccessible. Scheduling was halted during the file system recovery.

June 28, 2022: Unscheduled outage – home file system migration
Due to the delay in the return to service of Theta and ThetaGPU after a scheduled maintenance, the file systems were not available to the users.

August 1, 2022: Unscheduled outage – cooling system
All ALCF systems were shut down due to a cooling system failure. A faulty hose connection resulted in one of the chillers shutting down. The remaining chiller did not have enough capacity and temperatures in the data center rose rapidly. ALCF quickly shut down all hardware to avoid a catastrophic hardware loss. Within two days, the cooling system was repaired, and ALCF restarted the systems.

2.6.1.2 MTTI and MTTF

MTTI and MTTF Summary

MTTI and MTTF are reportable values with no specific targets. Table 2.14 summarizes the current MTTI and MTTF values.

Table 2.14 MTTI and MTTF Results

Theta-fs0 File System Cray Sonexion 2000 with 9.2 PB of usable storage				
	CY 2021		CY 2022	
	Target	Actual	Target	Actual
System MTTI	N/A [a]	15.25 days	N/A	13.52 days
System MTTF	N/A	182.47 days	N/A	72.51 days

[a] N/A = not applicable.

Theta-fs0 currently follows the Theta biweekly maintenance schedule. Theta-fs0 is not necessarily unavailable when Theta is in PM, but the PMs are often used to apply upgrades and patches.

2.6.2 Grand and Eagle Lustre File Systems

The ALCF installed a new set of Lustre file systems in 2021 running a Cray E1000 storage solution. Grand and Eagle each offer 100 PB of storage at 650 GB/s bandwidth and provide availability protection if one fails. Additionally, the file systems have the capability of sharing via Globus, a move toward providing a community file system. The file systems went into production on January 1, 2022.

ALCF proposed target metrics of 90 percent overall availability and 90 percent scheduled availability since these file systems are tightly integrated with Theta, which has the same target metrics.

2.6.2.1 Grand Scheduled and Overall Availability

Table 2.15 summarizes the availability results for the Grand file system.

Table 2.15 Availability Results – Grand

Grand File System Cray E1000 with 100 PB of storage at 650 GB/s				
	CY 2021		CY 2022	
	Target (%)	Actual (%)	Target (%)	Actual (%)
Scheduled Availability	N/A [a]	100.0	90	99.5
Overall Availability	N/A	96.5	90	96.6

[a] N/A = not applicable.

The remainder of this section covers significant availability losses. Details on the calculations can be found in Appendix A.

Grand – Explanation of Significant Availability Losses

This section briefly describes the causes of major losses of availability of Grand for the period of January 1, 2022, through December 31, 2022, as noted in Figure 2.12.

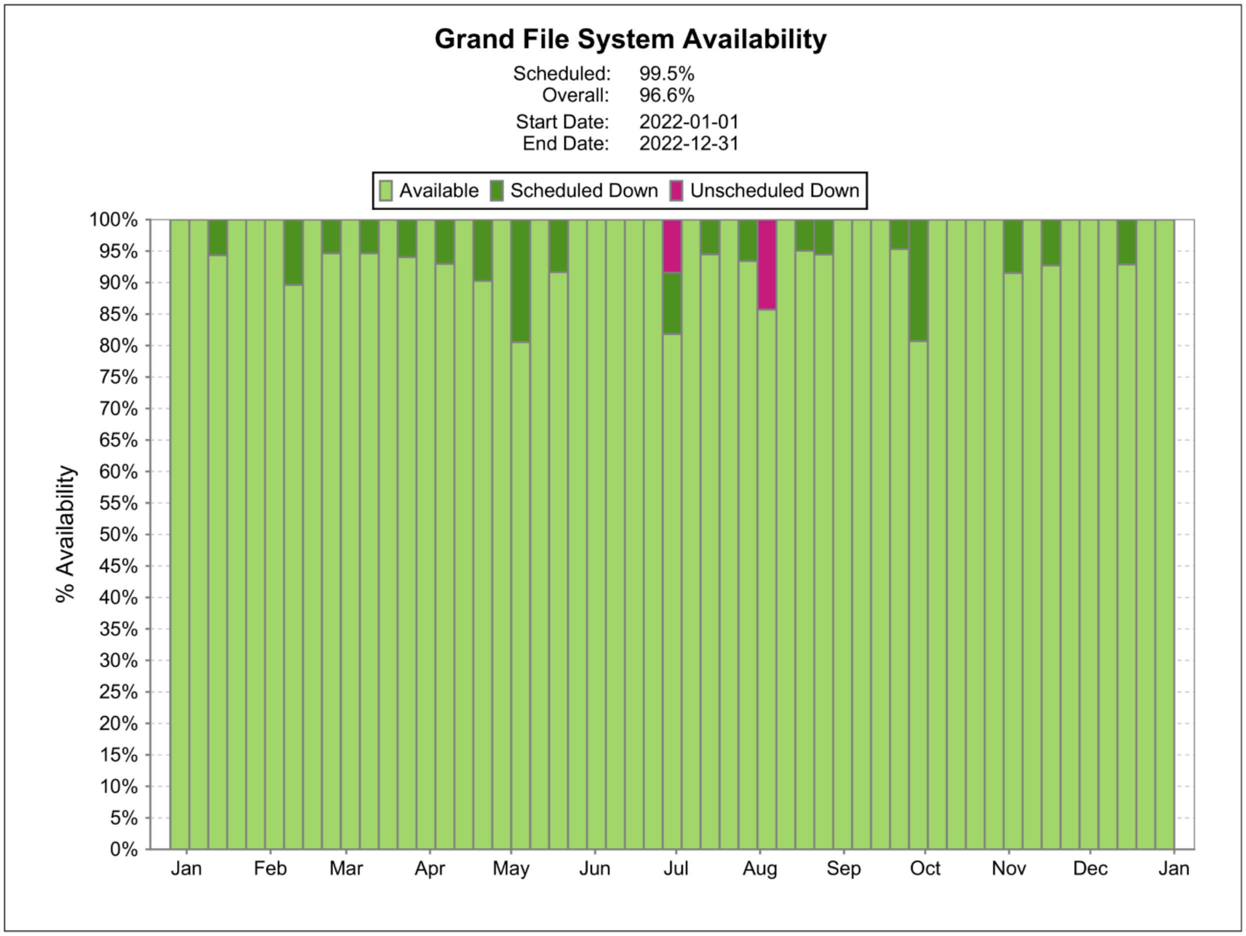

Figure 2.12 Grand Weekly Availability for CY 2022

Graph Description: Each bar in Figure 2.12 represents the percentage of the machine available for seven days. Each bar accounts for time in one of three categories. The pale-green portion represents available node-hours; the darker green represents scheduled downtime for that week; and magenta represents unscheduled downtime. Each of the significant loss events is described in detail below (these also appeared in the Theta section).

June 28, 2022: Unscheduled outage – home file system migration
Due to the delay in the return to service of Theta and ThetaGPU after a scheduled maintenance, the file systems were not available to the users.

August 1, 2022: Unscheduled outage – cooling system
All ALCF systems were shut down due to a cooling system failure. A faulty hose connection resulted in one of the chillers shutting down. The remaining chiller did not have enough capacity and temperatures in the data center rose rapidly. ALCF quickly shut down all hardware to avoid a catastrophic hardware loss. Within two days, the cooling system was repaired, and ALCF restarted the systems.

Grand – MTTI and MTTF Summary

MTTI and MTTF are reportable values with no specific targets. Table 2.16 summarizes the current MTTI and MTTF values.

Table 2.16 MTTI and MTTF Results – Grand

Grand File System Cray E1000 with 100 PB of storage at 650 GB/s				
	CY 2021		CY 2022	
	Target	**Actual**	**Target**	**Actual**
System MTTI	N/A [a]	15.31 days	N/A	14.68 days
System MTTF	N/A	182.45 days	N/A	121.13 days

[a] N/A = not applicable.

Grand generally follows Theta’s biweekly maintenance schedule. Grand is not necessarily unavailable when Theta is in maintenance, but the maintenance windows are often used to apply upgrades and patches.

2.6.2.2 Eagle – Scheduled and Overall Availability

Table 2.17 summarizes the availability results for the Eagle file system.

Table 2.17 Availability Results – Eagle

Eagle File System Cray E1000 with 100 PB of storage at 650 GB/s				
	CY 2021		CY 2022	CY 2022
	Target (%)	Actual (%)	Target (%)	Actual (%)
Scheduled Availability	N/A [a]	100.0	N/A	99.5
Overall Availability	N/A	96.5	N/A	96.6

[a] N/A = not applicable.

The remainder of this section covers significant availability losses. Details on the calculations can be found in Appendix A.

Eagle – Explanation of Significant Availability Losses

This section briefly describes the causes of major losses of availability of Eagle for the period of January 1, 2022, through December 31, 2022, as annotated in Figure 2.13.

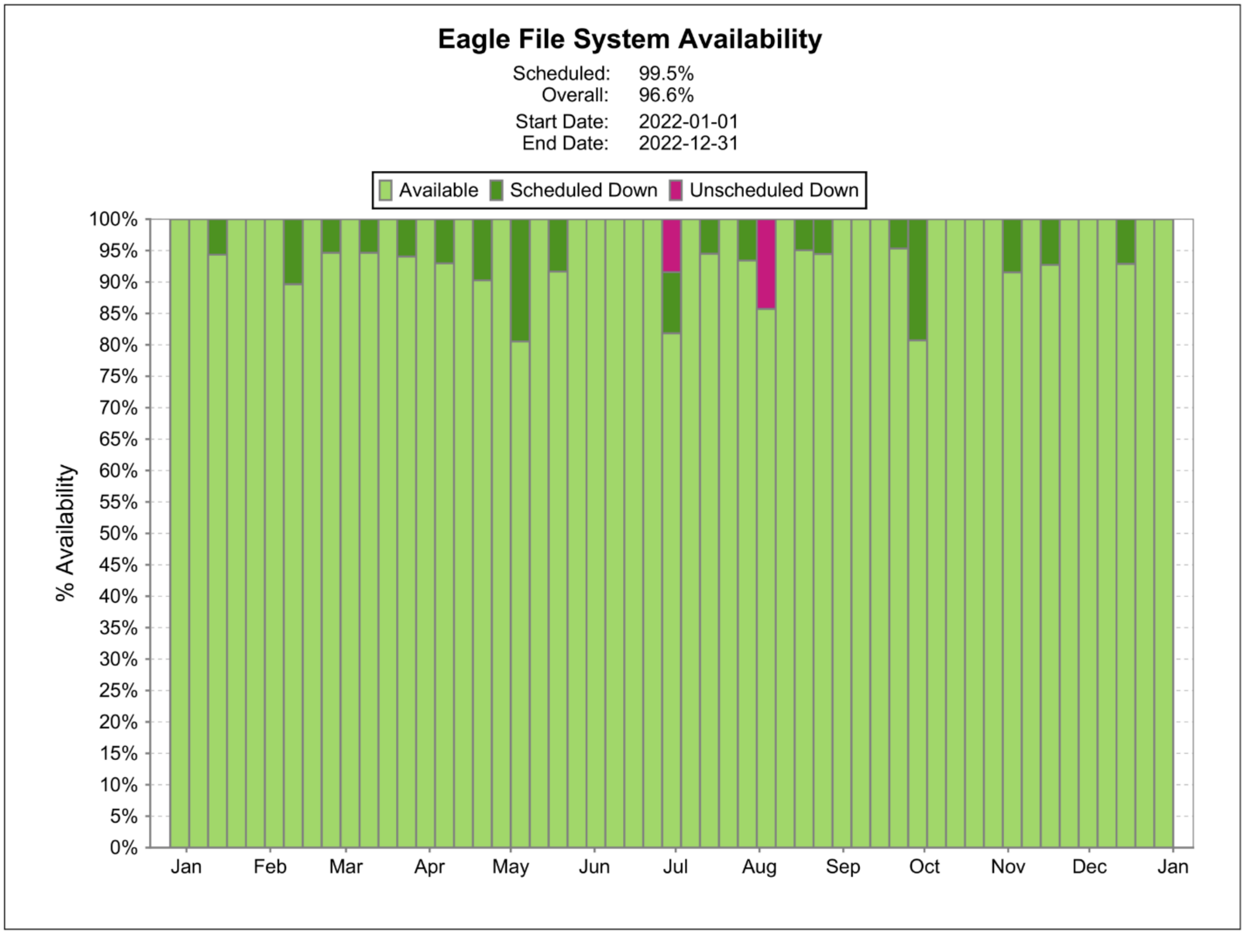

Figure 2.13 Eagle Weekly Availability for CY 2022

Graph Description: Each bar in Figure 2.13 represents the percentage of the machine available for seven days. Each bar accounts for time in one of three categories. The pale-green portion represents available node-hours; the darker green represents scheduled downtime for that week; and magenta represents unscheduled downtime. Each of the significant loss events is described in detail below (these also appeared in the Theta section).

June 28, 2022: Unscheduled outage – home file system migration
Due to the delay in the return to service of Theta and ThetaGPU after a scheduled maintenance, the file systems were not available to the users.

August 1, 2022: Unscheduled outage – cooling system
All ALCF systems were shut down due to a cooling system failure. A faulty hose connection resulted in one of the chillers shutting down. The remaining chiller did not have enough capacity and temperatures in the data center rose rapidly. ALCF quickly shut down all hardware to avoid a catastrophic hardware loss. Within two days, the cooling system was repaired, and ALCF restarted the systems.

Eagle – MTTI and MTTF Summary

MTTI and MTTF are reportable values with no specific targets. Table 2.18 summarizes the current MTTI and MTTF values.

Table 2.18 MTTI and MTTF Results – Eagle

Eagle File System Cray E1000 with 100 PB of storage at 650 GB/s				
	CY 2021		CY 2022	
	Target	Actual	Target	Actual
System MTTI	N/A [a]	16.01 days	N/A	14.68 days
System MTTF	N/A	365.00 days	N/A	121.13 days

[a] N/A = not applicable.

Eagle generally follows Theta's biweekly maintenance schedule. Eagle is not necessarily unavailable when Theta is in maintenance, but the maintenance windows are often used to apply upgrades and patches.

2.6.3 Tape Storage

The facility-wide high-performance storage system (HPSS) tape archive was available to all ALCF users from all compute resources in 2022, as in previous years. The tape storage is comprised of three 10,000-slot libraries with LTO8 tape drives and LTO8 tapes, with some legacy LTO6 drives and tapes remaining. The first tape library went into production in 2009 in the old Interim Supercomputing Support Facility (ISSF) datacenter, and the second followed in 2010 in the TCS datacenter. The third library went into production in 2016. In 2019, all of the tape libraries were moved to another building to provide separation of the archive data from the data center while also permanently vacating the ISSF datacenter. The HPSS disk cache and data movers are in the TCS datacenter. With the LTO8 drives and tape technology, the tape libraries have a maximum storage capacity of 305 PB.

2.6.3.1 Scheduled and Overall Availability

ALCF uses the target metrics of 90 percent overall availability and scheduled availability as proposed in the OAR for CY 2021. Table 2.19 summarizes the availability results.

Table 2.19 Availability Results

HPSS Archive LTO8 tape drives and tape with 350 PB storage capacity				
	CY 2021		CY 2022	
	Target (%)	Actual (%)	Target (%)	Actual (%)
Scheduled Availability	N/A [a]	100.0	90	99.7
Overall Availability	N/A	96.5	90	96.7

[a] N/A = not applicable.

Note that HPSS is considered unavailable when users can’t retrieve or access files via logins or data transfer nodes even though the HPSS libraries were unaffected during the scheduled maintenance periods, and still could do system functions like data migration. Therefore, HPSS overall availability reflects that users could not access it during scheduled maintenance.

Explanation of Significant Availability Losses

This section briefly describes the causes of major losses of availability of HPSS for the period of January 1, 2022, through December 31, 2022, as annotated in Figure 2.14.

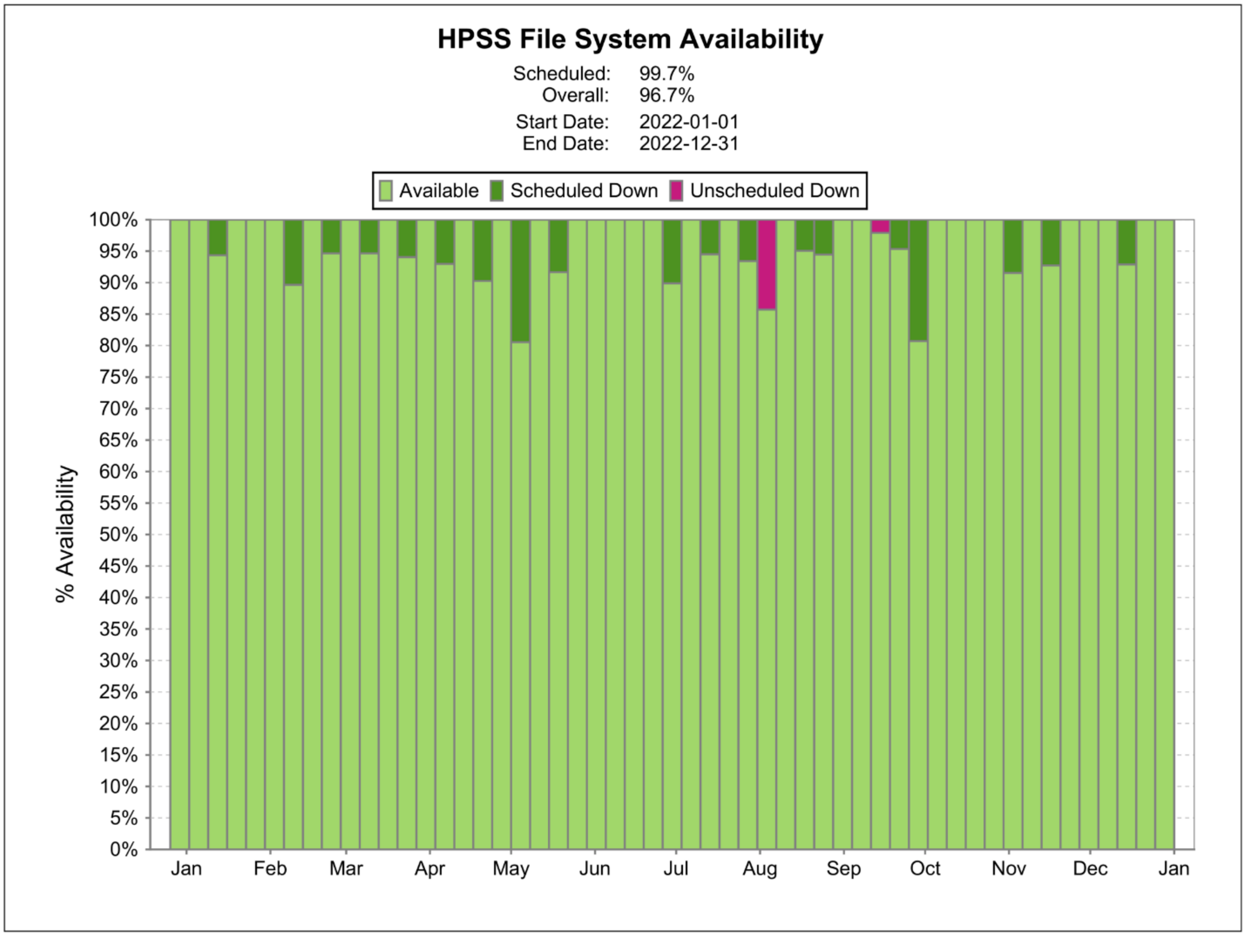

Figure 2.14 HPSS Weekly Availability for CY 2022

Graph Description: Each bar in Figure 2.14 represents the percentage of the file system availability for seven days. Each bar accounts for time in one of three categories. The pale-green portion represents percent-availability; the darker green represents scheduled downtime of the machines for that week; and magenta represents unscheduled downtime. Each of the significant loss events is described in detail below.

August 1, 2022: Unscheduled outage – cooling system
All ALCF systems were shut down due to a cooling system failure. A faulty hose connection resulted in one of the chillers shutting down. The remaining chiller did not have enough capacity and temperatures in the data center rose rapidly. ALCF quickly shut down all hardware to avoid a catastrophic hardware loss. Within two days, the cooling system was repaired, and ALCF restarted the systems.

September 9, 2022: Unscheduled outage – database failure
The database system for HPSS failed. Resolution: restarting the database system.

2.6.3.2 MTTI and MTTF

MTTI and MTTF Summary

MTTI and MTTF are reportable values with no specific targets. Table 2.20 summarizes the current MTTI and MTTF values.

Table 2.20 MTTI and MTTF Results

HPSS Archive LTO8 tape drives and tape with 350 PB storage capacity				
	CY 2021		CY 2022	
	Target	**Actual**	**Target**	**Actual**
System MTTI	N/A [a]	16.01 days	N/A	14.70 days
System MTTF	N/A	365.00 days	N/A	121.28 days

[a] N/A = not applicable.

HPSS maintenance is not regular but typically aligned with Theta’s maintenance schedule. HPSS is often available even though other resources may be in preventative maintenance.

2.7 Center-Wide Operational Highlights

2.7.1 Polaris

ALCF leveraged the Jenkins and ReFrame open-source frameworks to perform the acceptance testing (AT) process for Polaris as well as acceptance of the Polaris upgrade. The tests executed scientific applications, which covered a range of domains and characteristics. These included HACC, QMCPack, Nek, LAMMPS, and CosmicTagger. Each of these scientific applications was run at various scales up to full machine jobs with the expected output and performance validated. In all, 70 different scientific application configurations were utilized for the Polaris AT process. ReFrame is a python framework for the development of HPC system tests and includes built-in support for HPC job schedulers, including Polaris’s job scheduler, PBS, allowing for the easy submission of jobs and tracking of job progress. To automate the execution of the ReFrame tests, ALCF employed the Jenkins continuous integration (CI) framework. Jenkins provides a convenient interface and hooks to enable effective test failure triage. Jenkins enabled ALCF to produce continual submissions of the ReFrame tests, submitting a new execution of a given test after the completion of the previous execution.

Polaris was the first production deployment of PBS for job submission, execution, and management. All future ALCF machines, when possible, will use PBS. The scheduler has the flexibility to allocate nodes and individual GPUs to batch jobs, interactive jobs, and on-demand jobs. ALCF developed python code using the hook feature that PBS executes for specific events such as job submission and completion. Other scripts work with PBS to check the health of the selected nodes prior to the start of the job.

ALCF has been working with the Advanced Photon Source (APS) at Argonne to provide on-demand computing using PBS. The goal is a fully automated workflow driven by the beamline that moves the data to ALCF, computes on it, puts the results on our Eagle file system, and makes it visible to the right people via Globus share so that they can use the results to adjust the experiment.

2.7.2 AI Testbed

The ALCF AI Testbed went into production on May 2, 2022. It houses some of the most advanced AI accelerators for scientific research.

The goal of the testbed is to enable explorations into next-generation machine learning applications and workloads, enabling the ALCF and its user community to help define the role of AI accelerators in scientific computing and how to best integrate such technologies with supercomputing resources.

The Cerebras CS-2 system is a wafer-scale deep learning accelerator comprising 850,000 processing cores, each providing 48 KB of dedicated SRAM memory for an on-chip total of 40 GB and interconnected to optimize bandwidth and latency. Its software platform integrates popular machine learning frameworks such as TensorFlow and PyTorch.

The SambaNova DataScale system is architected around the next-generation Reconfigurable Dataflow Unit (RDU) processor for optimal dataflow processing and acceleration. The AI Testbed’s SambaNova system is a half-rack system consisting of two nodes, each of which features eight RDUs interconnected to enable model and data parallelism. SambaFlow, its software stack, extracts, optimizes, and maps dataflow graphs to the RDUs from standard machine learning frameworks like PyTorch.

The SambaNova system was upgraded to the SambaNova Gen2 system. An upgrade of the Cerebras system is expected in early 2023. The new installation is leveraging Kubernetes internally and is under active development by the vendor.

2.7.3 Sunspot / Aurora

ALCF installed Sunspot, Aurora’s Test and Development System (TDS), which consists of 2 racks, each with 64 nodes, for a total of 128 nodes. Each node consists of two Intel Xeon CPU Max Series (codename Sapphire Rapids or SPR) and six Intel Data Center GPU Max Series (codename Ponte Vecchio or PVC). Each Xeon has 52 physical cores supporting 2 hardware threads per core. The interconnect is HPE Slingshot 11, and PBS is the job scheduler. ALCF opened Sunspot to ECP and ESP users in December of 2022.

Built by Intel and Hewlett Packard Enterprise (HPE), Aurora will be theoretically capable of delivering more than two exaflops of computing power. Over the past year, ALCF installed computer racks and several components including the Intel DAOS (Distributed Asynchronous Object Storage) system and Intel’s Ponte Vecchio GPUs and Sapphire Rapids CPUs.

ALCF has granted Sunspot access to a limited number of users from the Early Science Program (ESP) and the Exascale Computing Program (ECP).

2.7.4 Storage Changes

ALCF migrated the GPFS mira-home filesystem to a new Lustre-based swift-home filesystem. The move was transparent to the users. The swift-home file system is accessible from Theta, ThetaGPU, Cooley, and Polaris.

ALCF performed live upgrades to the storage hardware for the Grand and Eagle storage systems. There were significant pauses in I/O during failovers; however, no jobs failed during the upgrades. When HPE discovered a bug in the ClusterStor version installed in September that caused filesystem corruption, ALCF did an emergency upgrade in late November and into early December. See section 5.5 for more details. ALCF was able to apply the upgrade with the filesystem still mounted and available by failing over high availability (HA) pair nodes, then failing back.

ALCF deployed Globus Connect V5, the latest Globus Transfer version. With this, ALCF is transitioning to having an endpoint per filesystem, rather than per-compute system. This allows the stopping of transfers on a filesystem-by-filesystem basis, rather than having to stop all transfers per compute system.

Conclusion

ALCF is maximizing the use of its HPC systems and other resources consistent with its mission. ALCF has exceeded the metrics of system availability and capability hours delivered. For the reportable areas—MTTI, MTTF, and utilization—ALCF is on par with OLCF and NERSC, and the values reported are reasonable. These measures are summarized in Table 2.1. ALCF exceeded the metrics for INCITE hours on Theta.

Storage system upgrades, especially live upgrades, enabled ALCF to maximize the availability of the ALCF compute resources.

Section 3. Allocation of Resources

(a) Did the allocation of computing resources conform with ASCR's published allocation policies (i.e., ratio of resources allocated between INCITE, ALCC, Director's Discretionary, ECP)?

(b) Was the allocation of Director's Discretionary computing resources reasonable and effective?

(c) Did the Facility encounter issues with under- or over-utilization of user allocations? If so, was the Facility's management of these issues effective in maximizing productive use of resources while promoting equity across the user population?

ALCF Response

The allocation of resources is consistent with ASCR's requested allocation policies. The breakdown of allocations available to INCITE, ALCC, DD, and ECP is 60 percent, 20 percent, 10 percent, and 10 percent, respectively. The 30 percent of the facility's resources available to ASCR is provided through the 20 percent to ALCC and 10 percent to ECP.

The INCITE program fully allocates the 60 percent of the time available to it. The DD time is overallocated, but substantially underused due to the exploratory nature of the projects in the DD program.

As the results show in Section 8, these are reasonable allocations of resources. Below are a few areas ALCF considers when analyzing usage statistics for the allocation programs.

3.1 Usage of the INCITE and ALCC Hours

The 2022 INCITE program allocated 17.8M node-hours to 17 projects on Theta (with expansion) and 1.3M node-hours to 10 projects on Polaris. The allocation usage by project is shown in Figure 3.1 for Theta (with expansion) and Figure 3.2 for Polaris. A total of 20.4M node-hours were delivered to INCITE projects on Theta (with expansion) (Table 3.1). Of these 17 projects, only 4 used less than 75 percent of their allocation. The other 13 projects used more than 87 percent of their allocations, including 8 projects that used their entire allocations or more. INCITE and ALCC projects were permitted to exceed their node-hour allocations up to 125 percent, enabled by the ALCF overburn policy that permitted projects to continue running capability-sized jobs after their allocations are exhausted (as explained in section 3.3).

A total of 589.6K node-hours were delivered to INCITE projects on Polaris (Table 3.2). Of these 10 projects, 5 used less than 10 percent of their allocation. The other 5 projects were able to use more than 48 percent of their allocations. Three of those projects used more than 70 percent of their allocation, including one that used more than 99 percent of their allocation. The primary challenges in the delivery of Polaris time were the lack of information when the INCITE call was open and Polaris was moving into production later than originally planned. Most projects given INCITE time on Polaris were exploratory projects with a primary allocation on another system.

When Polaris became available, many PIs did not prioritize that exploration, as it was deep into the INCITE year.

For the 2021–2022 ALCC year, 18 projects had allocations on Theta (with expansion) with a total of 6.8M node-hours used. The usage is shown in Figure 3.3. Two of the projects used less than 1 percent of their allocation, 5 projects used between 80 and 100 percent, and 11 projects used their entire allocation or more. Polaris was not available for the 2021–2022 ALCC year.

The 2022–2023 ALCC was approximately halfway through its allocation cycle at the end of 2022. At that point, 11 projects received allocations totaling 5.9M node-hours on Theta (with expansion), and 6 projects received allocations totaling 800.0K node-hours on Polaris. The 2022–2023 ALCC projects on Theta (with expansion) used a total of 2.7M node-hours from July 1, 2022, through December 31, 2022, and their allocation usage is shown in Figure 3.4. The 2022–2023 ALCC projects on Polaris used a total of 92.6K node-hours from July 1, 2022, through December 31, 2022, and their allocation usage is shown in Figure 3.5.

Table 3.3 summarizes the ALCC node-hours allocated and used on Theta (with expansion) in CY 2022. This includes COVID-19 research projects run on Theta (with expansion) at the request of ASCR, which typically had off-cycle start and end dates. The 6.4M ALCC node-hours allocated are calculated by adjusting the 2021–2022 and 2022–2023 ALCC year allocations on Theta (with expansion) by the percentage of their award cycle occurring in CY 2022, then summing these two values. The total 7.4M ALCC node-hours used is the sum of all node-hours used by any ALCC project on Theta (with expansion) in CY 2022. Table 3.4 summarizes the ALCC node-hours allocated and used on Polaris in CY 2022. The 355.8K ALCC node-hours allocated are calculated by adjusting the 2022–2023 Polaris ALCC year allocations by the percentage of their award cycle occurring in CY 2022. The total 92.6K ALCC node-hours used is the sum of all node-hours used by any Polaris ALCC project in CY 2022.

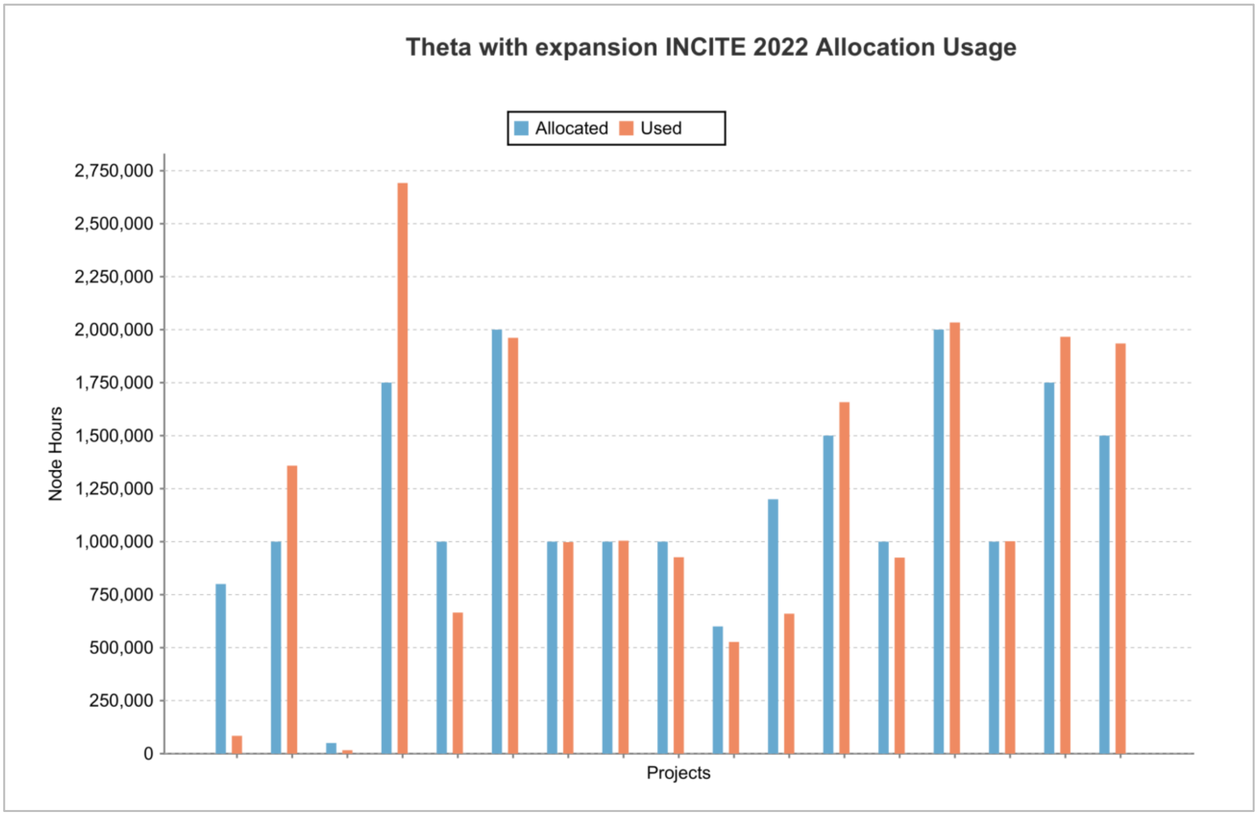

Figure 3.1 Theta (with expansion) INCITE 2022 Allocation Usage (Note: Projects are randomly ordered.)

Table 3.1 INCITE 2022 Time Allocated and Used on Theta (with expansion) in CY 2022

Projects	Theta (with expansion)
Allocated Node-Hours	17.8M
Used Node-Hours	20.4M

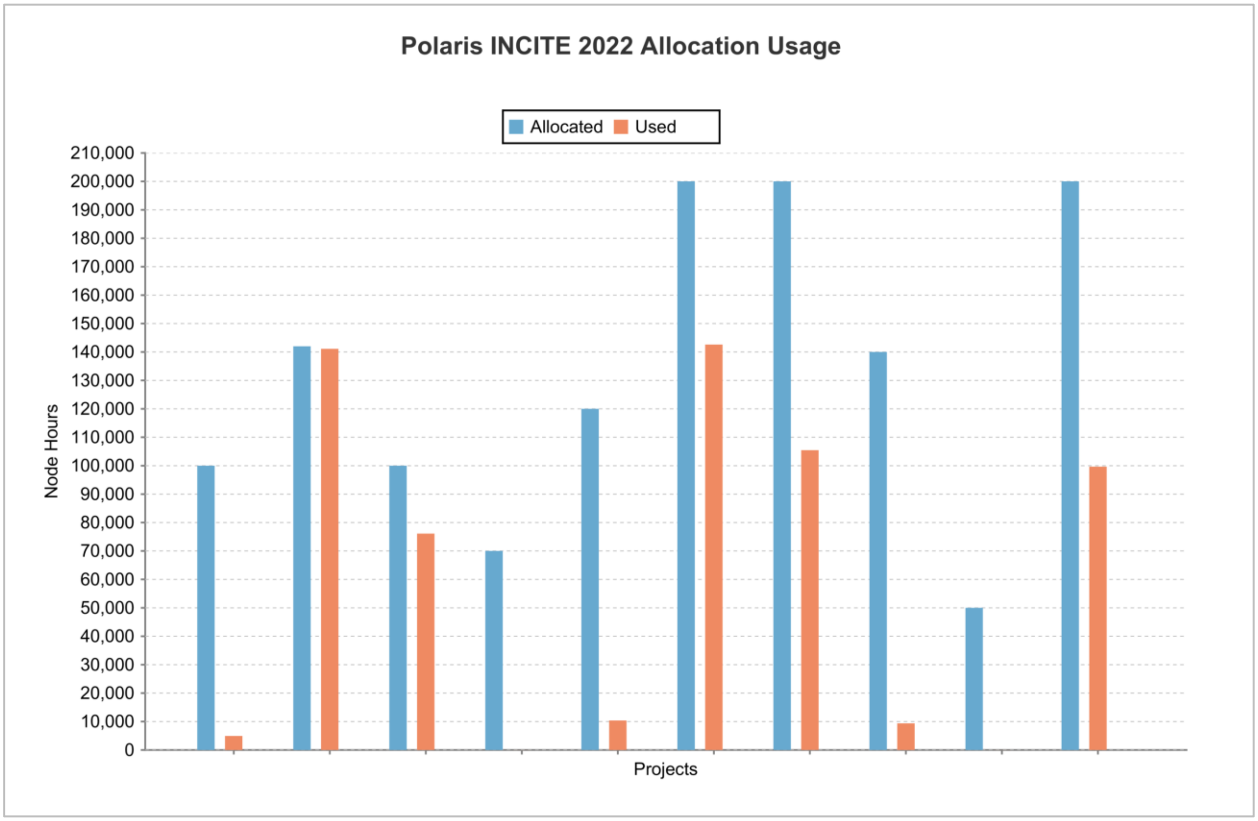

Figure 3.2 Polaris INCITE 2022 Allocation Usage (Note: Projects are randomly ordered.)

Table 3.2 INCITE 2022 Time Allocated and Used on Polaris in CY 2022

Projects	Polaris
Allocated Node-Hours	1322.0K
Used Node-Hours	589.6K

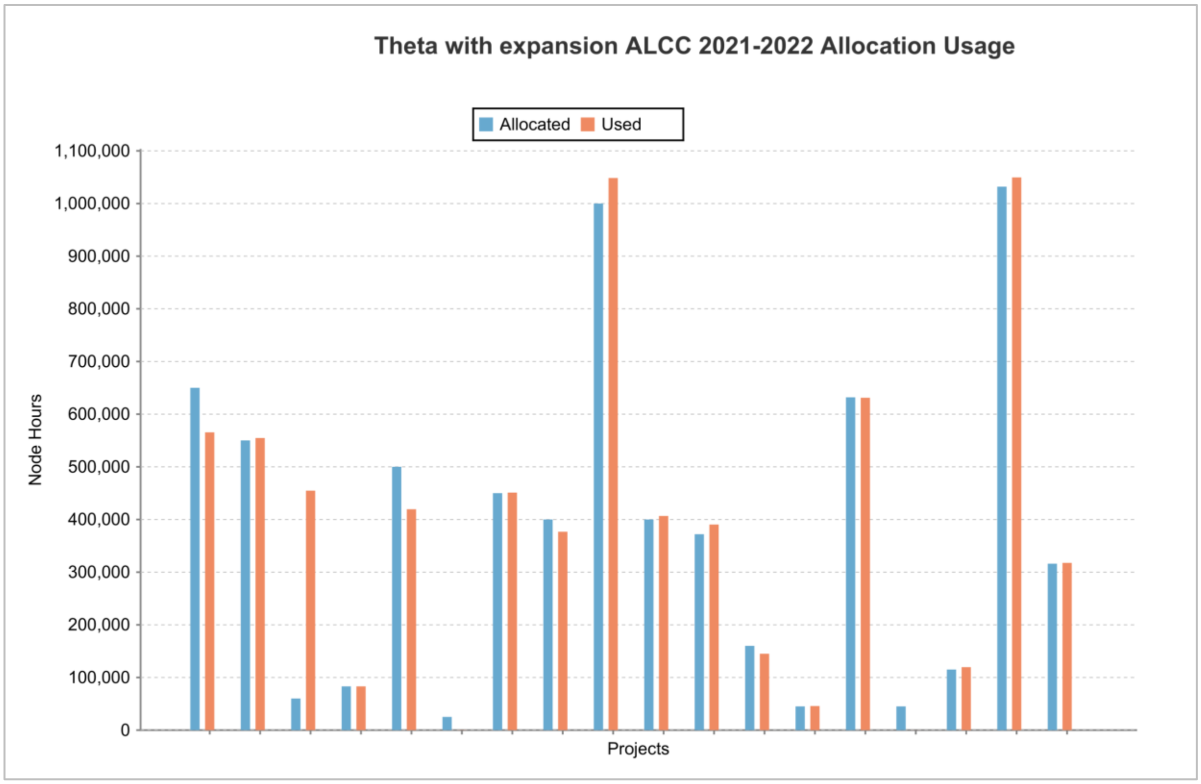

Figure 3.3 Theta (with expansion) ALCC 2021–2022 Allocation Usage (Note: Projects are randomly ordered.)

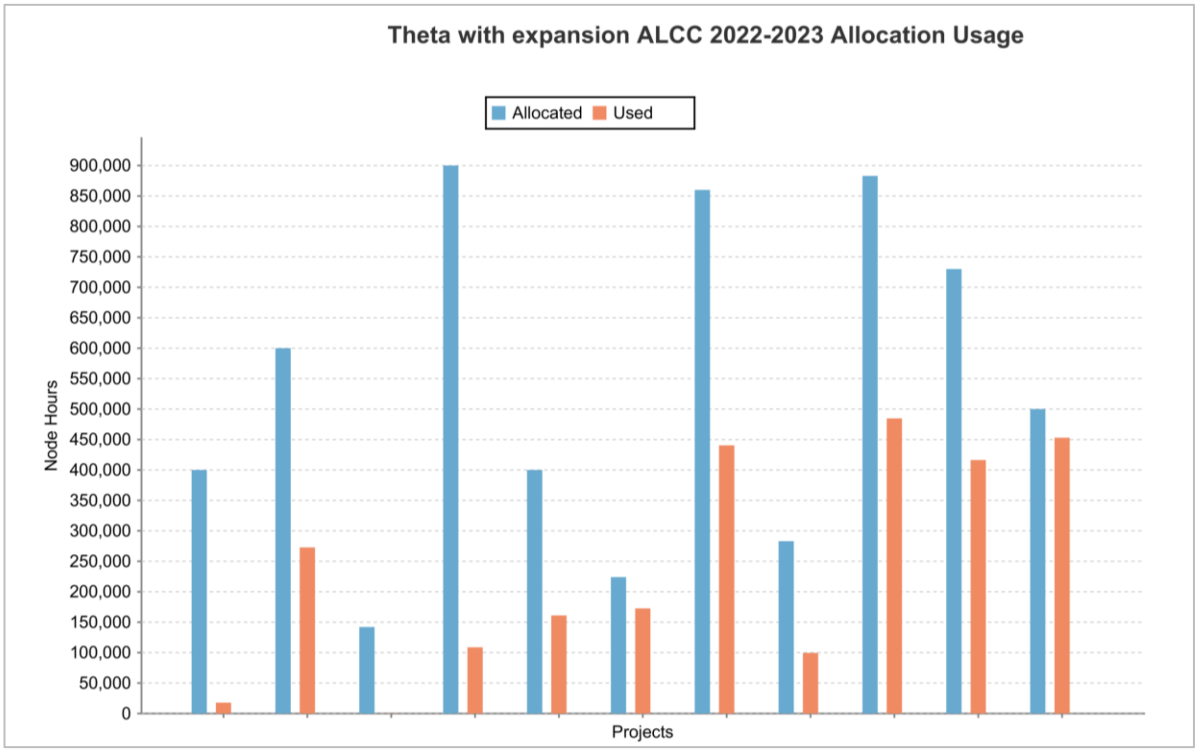

Figure 3.4 Theta (with expansion) ALCC 2022–2023 Allocation Usage as of December 31, 2022 (Note: Projects are randomly ordered.)

Table 3.3 ALCC Time Allocated and Used on Theta (with expansion) in CY 2022

Projects	Theta (with expansion)
Allocated Node-Hours	6.4M [a]
Used Node-Hours	7.4M [b]

[a] Allocation total is obtained by adjusting each of the ALCC cycle allocations (2021–2022, 2022–2023) to prorate for the amount of time allocated in CY 2022, then summing.

[b] Usage total is the number of node-hours charged for jobs run against any ALCC allocation in CY 2022.

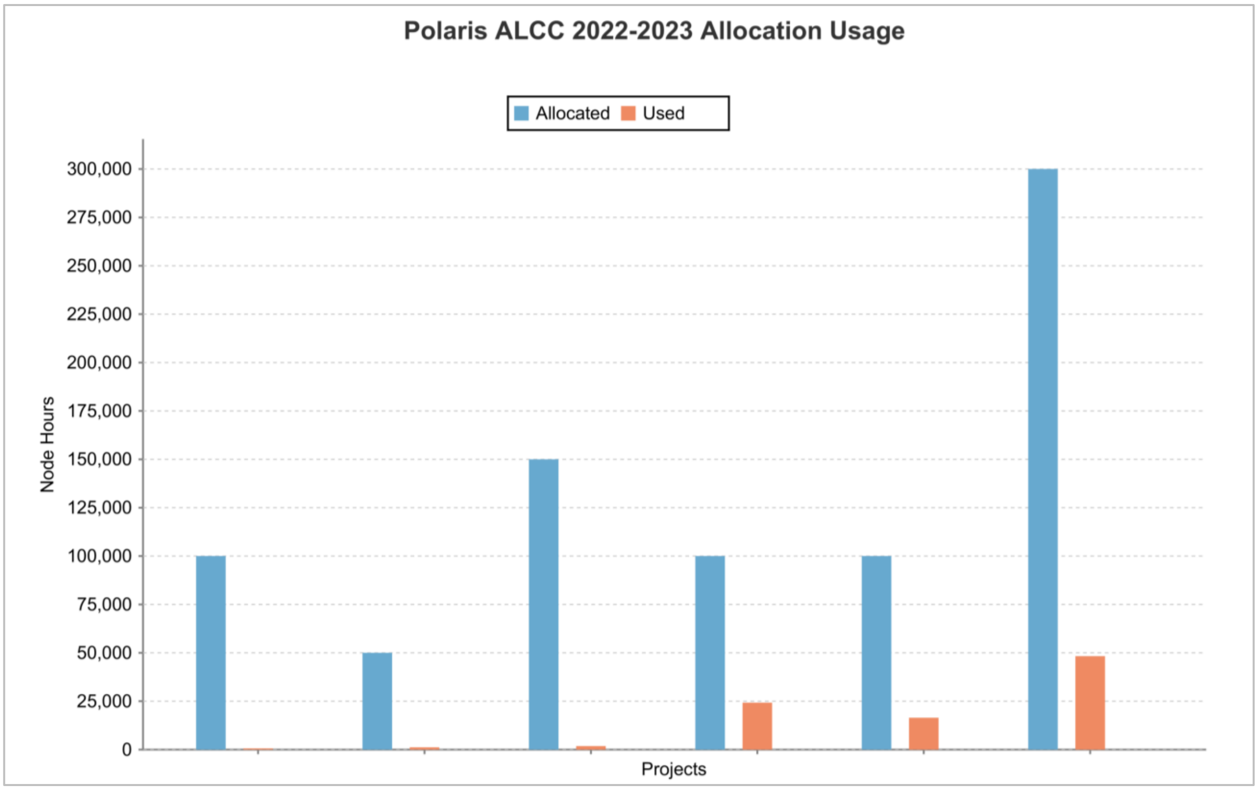

Figure 3.5 Polaris ALCC 2022–2023 Allocation Usage as of December 31, 2022 (Note: Projects are randomly ordered.)

Table 3.4 ALCC Time Allocated and Used on Polaris in CY 2022

Projects	Polaris
Allocated Node-Hours	355.8K [a]
Used Node-Hours	92.6K [b]

[a] Allocation total is obtained by adjusting the 2022–2023 ALCC cycle allocations to prorate for the amount of time allocated in CY 2022.

[b] Usage total is the number of node-hours charged for jobs run against any ALCC allocation in CY 2022.

3.2 Facility Director's Discretionary Reserve Time

The 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 science
4) Internal/support
5) ECP support

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

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

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

Internal/support projects are devoted to supporting the ALCF mission. ALCF does not reserve node-hours for internal activities. All internal use comes 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.

To support the dynamic needs of ECP, the ECP time was moved from ALCC to DD starting in 2019, but the 10 percent allocation is still part of the overall ASCR fraction of the system. As a result, the discretionary pool grew to 20 percent of the system to support ECP. ECP and the computing facilities run a Resource Allocations Council (RAC) that meets monthly to discuss the computing needs of ECP projects, allocating up to 10 percent of the system.

DD 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 program has high demand, and often the requested amount cannot be accommodated.

Tables 3.5 and 3.6 show the total time allocated and used in the DD program on Theta (with expansion) and Polaris, respectively, during 2022. By its very nature, the DD program is amenable to over-allocation because time is often left unused; however, it should be noted that these totals do not represent available allocations for the entire calendar year. A project might have a 1,000-node-hour allocation that only persists for three months, but that 1,000-node-hour allocation is counted entirely in the annual total node-hour number. Projects are not guaranteed the time allocated. DD projects run at a lower priority than INCITE or ALCC projects, which could reduce 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 (with expansion) in CY 2022

Projects	Theta (with expansion)
Allocated Node-Hours	7.1M
Used Node-Hours	5.3M

Table 3.6 DD Time Allocated and Used on Polaris in CY 2022

Projects	Polaris
Allocated Node-Hours	543.6K
Used Node-Hours	481.5K

Lists of the CY 2022 DD projects on Theta (with expansion) and Polaris are provided in Appendix B.

Figures 3.6 and 3.7 provide breakdowns of the CY 2022 DD allocations by domain for Theta (with expansion) and Polaris, respectively.

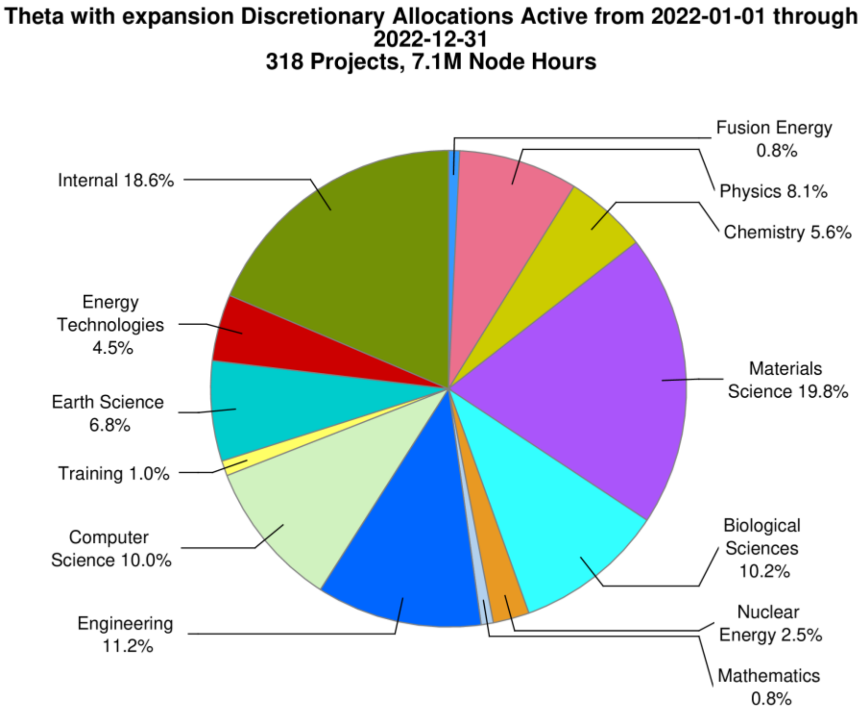

Figure 3.6 Theta (with expansion) CY 2022 DD Allocations by Domain

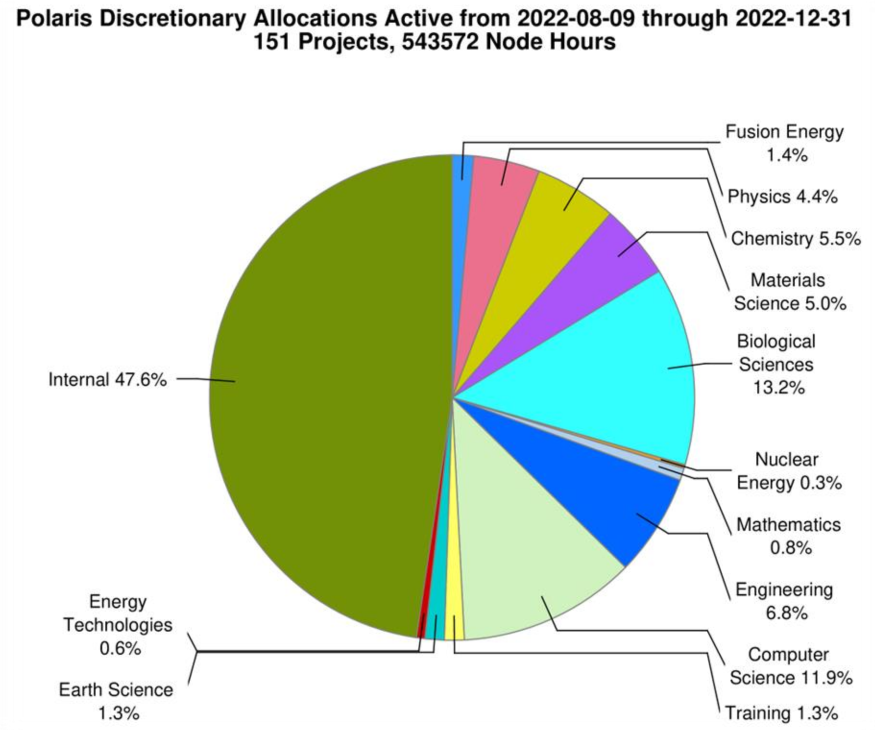

Figure 3.7 Polaris CY 2022 DD Allocations by Domain

3.3 User Allocation Utilization

Inevitably, some projects will use less time than allocated and other projects will want to use more. To rebalance some of the allocated time across projects to ensure optimal utilization of resources, an overburn policy is in effect for INCITE and ALCC projects, which permits high utilization projects to continue using the machine effectively for capability jobs. If an INCITE or ALCC project has exhausted its allocation in the first 11 months of its allocation year, it is eligible for overburn time. At this point, capability jobs submitted by INCITE and ALCC projects will run in the default queue (instead of backfill) for the first 11 months of the allocation year until 125 percent of the project allocation has been consumed.

Should additional overburn hours be needed, INCITE and ALCC projects may provide the facility with a short description of what the project plans to do with the additional hours, highlighting specific goals or milestones and the time needed to accomplish them. These requests are reviewed by the scheduling committee, allocations committee, and ALCF management. Non-capability jobs from projects that have exhausted their allocation will run in the backfill queue.

This overburn policy does not constitute a guarantee of extra time, and the facility reserves the right to prioritize jobs submitted by projects that have not yet used 100 percent of their allocation., The earlier that an INCITE or ALCC project exhausts its allocation, the more likely they are to be able to take full advantage of the overburn policy.

The ALCF has multiple methods of managing under- and over-utilization of user allocations. The overall goal of all the policies is to ensure that user projects have the greatest chance to accomplish their science goals.

Under-utilization earlier in the allocation year is primarily managed through personal contact with the projects to understand issues and assist with resolving any problems. Examples of these challenges can include data movement, scheduling challenges, porting problems and bugs. Most of the time this approach is effective, but if significant under-utilization persists farther into the allocation year, then ALCF may apply its stated policies for pulling back time from INCITE or ALCC projects to be redistributed among other projects in that program. For INCITE, utilization is assessed, and pullback can be applied at the beginning of the fifth and ninth months of the allocation year. For ALCC, utilization assessment and possible pullback occurs after the seventh month of the allocation year.

3.4 Other Large-Scale Managed Resources

In addition to its HPC systems, ALCF supports its user base with the Grand and Eagle file systems, which are accessible from multiple ALCF systems and constitute a vital part of the ALCF Community Data Co-Op. Every new project awarded compute time is also given a storage allocation on Grand or Eagle. Supplementary information on the technical specifications and availability for Grand and Eagle can be found in section 2.6.2, and a description of upgrades to these storage systems undertaken in 2022 can be found in section 2.7.4. As part of the mission to support collaborative and data-driven scientific discoveries, ALCF is providing researchers with services to securely share and reliably transfer data using Globus. In 2022, ALCF supported 36 data-only projects using these services. The Eagle filesystem with Globus sharing is a key component in ongoing efforts to automate the process of taking data at APS, processing that data with on-demand computing resources on Polaris, and finally sharing that processed data with collaborators and the scientific community as described in section 2.7.1.

Conclusion

The ALCF delivered the following node-hours to the allocation programs in CY 2022: 21.0 million to INCITE, 7.5 million to ALCC, and 5.8 million to DD. The DD program 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 jobs efficiently and to achieve science goals that could not otherwise have been reached.

This page intentionally left blank.

Section 4. Operational Innovation

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

(b) Have management/workforce innovations been implemented that have improved the facility's operations?

(c) Is the facility effectively utilizing their postdoctoral fellows?

ALCF Response

Listed below are the innovations and best practices carried out at ALCF during CY 2022. 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 facilities and beyond. As outlined in section 4.3, ALCF is effectively utilizing its postdoctoral fellows.

4.1 Operational Innovation – Technical

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

4.1.1 Working to Enable On-Demand Workloads

Challenge: The increasing need for on-demand analysis of instrument data, the results of which can serve as input into the next data acquisition, directly conflicts with the scheduling requirements of batch-oriented traditional HPC workloads.

Approach: Setting an attribute in PBS to designate a node as an "on-demand" node will automatically move it from the normal production queue to the on-demand queue. These nodes are also accessible via the preemptable queue where users are encouraged to do high-throughput computing (HTC)-type workloads (large numbers of relatively small, short jobs). This helps keep the nodes busy when on-demand jobs don't need the nodes, but the jobs can be preempted with minimal wasted computation to make room when on-demand jobs come in.

Impact/Status: ALCF has been extensively testing this approach on a small, 4-node testbed and has recently completed the first end-to-end tests on Polaris working with the APS. This innovation has not yet been shared with the other ASCR facilities.

4.1.2 Increased Availability of Logs Using Kafka

Challenge: HPE Performance Cluster Manager (HPCM) processes some, but not all, logs through Kafka, a system that is also critical to machine operations, potentially overloading Kafka with additional streams. This should be avoided, and an alternative solution should be pursued.

Approach: ALCF set up an additional Kafka instance, independent of machine operations, that uses xFilebeat from Elasticsearch to gather logs of interest that HPCM does not gather—

primarily, the PBS “mom” logs on the compute nodes—and then aggregates both Kafka streams in ALCF’s data warehouse. From there, ALCF can provide various views/representations of this data for interested parties. The first focus of this effort has been a per-job aggregation of all job-related log lines to help admins or users debug job failures.

Impact/Status: ALCF is currently testing this approach and plans to make the raw logs available to admins soon. Future work includes additional filtering and/or interpretation of the contents to make it more useful to users. This innovation has not yet been shared with the other ASCR facilities.

4.1.3 Making Filesystems a Scheduler-Aware Resource

Challenge: If the scheduler does not see the filesystems as a schedulable resource and cannot identify which jobs need which filesystems, scheduling must be stopped whenever a filesystem goes down, or jobs may fail when they attempt to perform I/O.

Approach: ALCF has added a custom resource to PBS to represent the filesystems and now requires users to specify which filesystems their jobs require at job submission time.

Impact/Status: This change, now in production, has allowed ALCF to continue running jobs during multi-day upgrades on production filesystems, because indicating which filesystem is down enables PBS to simply stop scheduling jobs that required that resource. This innovation has not yet been shared with the other ASCR facilities.

4.1.4 HIP on Aurora

Challenge: With two DOE machines, Frontier and El Capitan, featuring AMD GPUs, AMD’s Heterogeneous Interface for Portability (HIP) is a likely programming model for applications and runtimes to target. This raises the question of how to port codes using HIP to Aurora, which does not have native HIP support.

Approach: Porting HIP code using the natively supported SYCL programming model is a possibility, but offering native HIP support may also be a viable alternative which would avoid increasing application complexity and maintenance cost. The HIP on Aurora ECP project (also known as HIPLZ, or HIP on Level Zero) demonstrated HIP on Intel GPUs with the HIPLZ prototype, a HIP library that targets the native, low-level programming model of Aurora: Intel Level-Zero, a new API that Intel is developing for its GPUs, including the Ponte Vecchio GPUs in Aurora.

Impact/Status: Changes required for HIP in the LLVM infrastructure have been upstreamed, and the new unified OpenCL and Level-Zero HIP backend, called CHIP-SPV, leverages unmodified Clang/LLVM 14 or 15 and integrates seamlessly with AMD’s refactored HIP infrastructure. The objective for next year is to port three real applications of interest to ECP on top of CHIP-SPV, as well as their dependencies, demonstrating HIP is a viable programming model for applications on Aurora. The project is a collaboration with NERSC and OLCF; one paper has been published on the work, with a second in development on the latest efforts.

4.1.5 Benchmarking Resources in AI Testbed

Challenge: HPC centers are evaluating emerging novel hardware AI accelerators to efficiently run AI-driven science applications. The ALCF AI Testbed deploys novel AI accelerators from SambaNova, Cerebras, Graphcore, Groq, and Habana to evaluate their efficacy in accelerating scientific machine learning applications. With a wide diversity in the hardware architectures and software stacks of these systems, it is challenging to understand how these accelerators perform. The state-of-the-art in the evaluation of DL (Deep Learning) workloads primarily focuses on CPUs and GPUs.

Approach: This AI Testbed benchmarking project is a first-of-its-kind evaluation of these accelerators with diverse workloads, done in close collaboration with the vendors. The workloads included DL primitives, standard MLPerf benchmark models, and scientific ML applications. ALCF also evaluated the performance of collective communication, which is key for distributed DL implementation, along with a study of scaling efficiency.

Impact/Status: This project produced several key insights and opportunities for integrating these novel AI accelerators with supercomputing systems. The findings were presented at the 13th IEEE International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS22) at SC22.

4.2 Operational Innovation – Management/Workforce

ALCF works to prepare for next-generation systems through collaboration with vendors and other DOE facilities. ALCF reports on participation in research projects in Section 1.3.2 and professional community activities in Section 8.3.1.

4.2.1 Minimizing Cost and Effort around Managing SSL Certificates

Challenge: Commercial SSL (Secure Sockets Layer) certificates can be expensive. InCommon provides SSL certificates that are free but limited to a 1-year lifetime.

Approach: ALCF's infrastructure team developed an automated system to update SSL certificates that are due to expire. This capability has been integrated with ALCF's internal Vault secrets management system and SaltStack configuration management system.

Impact/Status: Certificates within 30 days of expiring are automatically renewed and moved to the new automated system. This results in less likelihood of certificates expiring (the system runs multiple times a day) and has eliminated five steps that the ALCF operations team had to perform when executing the process manually.

4.2.2 ALCC and INCITE Project Introduction and Welcome Calls

Challenge: It can be challenging to properly assess the readiness of new ALCC and INCITE project teams before the projects have begun. It is critical for ALCF staff and project teams to remain engaged early in a project to help ensure project teams' success on ALCF resources and overcome any difficulties in using their allocations.

Approach: In prior years, ALCF points-of-contact (e.g., Catalysts) held formal introductory welcome calls with ALCC and INCITE project teams in the month before the start of new project allocations to meet the principal investigator (PI) and team, review facility policies, provide helpful tips, and discuss where ALCF assistance might be most beneficial. This year, ALCF points-of-contact started initial email discussions with the PIs and team members about one month before their allocations started, sometimes augmented with phone and virtual discussions, and continued as teams began their work on ALCF resources. The formal welcome calls took place before the end of the first month, when team members are more likely to have questions and issues to discuss based on their experience to date, thus providing a stronger opportunity to engage with staff.

Impact/Status: Initial email discussions provided an early opportunity to help clear up questions that the PIs and their teams had as they prepared to use their allocations effectively on day one. Starting the discussions earlier and conducting the formal welcome calls later gave staff and the teams an opportunity to discuss any issues encountered early on. Examples from Polaris discussions included guidance for linking NVIDIA libraries and providing a workaround to an MPI-GPU issue while it investigated. Interactions like this help teams to make progress early and improve communication between staff and the project teams.

4.2.3 AI-Related Training

Challenge: HPC can be daunting and is not well promoted or taught to undergraduates. ALCF has developed resources and training materials to reach out to undergraduates in the U.S.

Approach: ALCF organized the second iteration of a training curriculum targeting undergraduates and graduate students, called "Introduction to AI-driven Science on Supercomputers." This eight-week course provided students with an introduction to AI/ML techniques applied to science problems and how to use the techniques on DOE supercomputers. Each week, the students had opportunities to apply AI/ML concepts in hands-on activities and heard from scientists at Argonne about how they were using those techniques in their work. The aim was to introduce supercomputing resources at an earlier stage of a scientist's career to broaden their toolkit.

Impact/Status: The 2022 program welcomed over 200 attendees from 90 universities, including undergraduates, graduate students, postdocs, and faculty. Of those attendees, 52 received a Certificate of Completion for the course, completing all eight homework assignments from each session. Feedback from attendees highlighted the value of sharing specific scientific problems that were solved throughout the series using the AI/ML tools.

4.3 Postdoctoral Fellows

ALCF supports a steady-state postdoctoral fellowship program. Within this program, ALCF supports one named postdoctoral fellow, the Margaret Butler for Computational Science Fellow. Postdocs are awarded one-year term with an option to renew for an additional year (this is typically the case), with a similar option for a third year. The major goal of the program is to either convert the postdocs to a regular staff position, place them at another DOE laboratory, or support their efforts to find an academic or industry position. The objective, in all cases, is for these postdocs to continue as lifelong users of DOE compute resources.

Applications for postdoctoral positions are handled by Argonne's Postdoctoral Program Office. In 2022, ALCF supported 16 postdoctoral researchers (fellows), including seven new ones, representing a range of scientific domains. During the year, none were terminated, and one transferred to a staff position in another Argonne division. Six postdocs were part of ESP projects, two were part of ECP projects and eight were part of ALCF's steady state operations. The ALCF postdocs worked on various research topics, including computational chemistry; computational fluid dynamics; regional climate modeling; numerical cosmology; heterogeneous frameworks; AI and Internet of Things (IoT) applications to environmental modeling; parallel computing; dark matter; performance engineering; data management, storage, and I/O; biomedicine; cosmology data analysis; quantum simulation of low dimensional systems; reduced-order modeling of high dimensional fluid dynamics systems; and mixed precision optimization for next-generation AI accelerator systems.

Once hired, each postdoc was assigned both a direct research supervisor and an Argonne staff mentor. The mentor, initially selected by the division or the supervisor, could be changed by the postdoc. The supervisors met with the postdoc on a weekly basis and also engaged in the postdoc's research efforts. The supervisor evaluated the progress and completed a yearly standardized review that was submitted to ALCF management consideration and authorization of appointment renewals. The mentor was responsible for meeting with the postdoc to discuss career development milestones and personal goals. This interaction happened as needed, but no less than once a quarter. The guidance for these discussions included key skills the postdoc should focus on over the next year; opportunities for development; and, if entering the third year, what skills or experience will be most beneficial to enabling a smooth career transition. The Division Director also met monthly with the postdocs as a group to hear progress updates, address any issues specific to the postdoc community, and solicit general feedback.

ALCF supported the following 16 postdoctoral researchers in CY 2022:

Riccardo Balin (Ph.D., aerospace engineering, University of Colorado Boulder). **Hired:** January 2021. **Research area:** Computational fluid dynamics (CFD), wall-bounded turbulence, data-driven and ML techniques for turbulence modeling, in situ coupling of simulation and AI/ML workloads. **Current projects:** (1) *CFDML_AESP*: Data Analytics and Machine Learning for Exascale CFD (primary project), and (2) *PHASTA_AESP*: Development of the PHASTA CFD Code for Exascale Simulations on Aurora. **Scientific goal:** To develop a framework for applying in-situ machine learning to subgrid modeling for hybrid Reynolds-averaged Navier-Stokes (RANS)/large eddy simulation (LES) of high Reynolds number separated flows. These coupled simulations will be among the first to run on Aurora. **Accomplishments:** (1) Developed a scalable infrastructure for in situ coupling of simulation and AI/ML with SmartSim at LCF. This infrastructure enables on-the-fly training of ML models from data produced by an ongoing simulation and the use of ML models within the CFD code PHASTA; (2) responsible for ensuring SmartSim tools build and run efficiently on LCF machines; (3) supported LCF's Data Science team efforts on Polaris to run scaling and performance benchmarks with various ML frameworks and distributed training libraries; (4) published an article in *Journal of Computational Physics* and produced a tutorial for ALCF's 2022 SDL workshop.

Shivam Barwey (Ph.D., aerospace engineering, University of Michigan). **Hired:** August 2022. **Research area:** Reduced-order modeling of high dimensional fluid dynamical systems, graph-based learning, data clustering, HPC. **Current project:** Interpretable data-based surrogate modeling for unsteady fluid flows in complex geometries using graph neural networks (GNNs). **Scientific goal:** To develop neural network-based surrogate models that (a) are compatible with unstructured/skewed meshes and geometric perturbations, and (b) produce interpretable latent spaces. **Accomplishments:** Current work is focused on the model development side (i.e., developing the interpretable GNN architecture, ensuring stability in predictions, compatibility with skewed meshes, etc.). Demonstrations are performed on data sourced from high-Re turbulent flow over backward-facing step configurations. Training is performed on Polaris nodes.

Denis Boyda (Ph.D., theoretical physics, National Research Center Kurchatov Institute, Moscow, Russia). **Hired:** September 2020. **Terminated:** January 2023: he is back at the Massachusetts Institute of Technology (MIT) as a fellow with the National Science Foundation (NSF) AI Institute for Artificial Intelligence and Fundamental Interactions (IAIFI). **Scientific goal:** To simulate the interactions of dark matter candidate particles and nuclei, leading to insights into dark matter and fundamental particle physics. **Accomplishments:** (1) Developed new sampling algorithms enforced with cutting-edge ML techniques for simulation of interesting physical phenomena such as the interaction of dark matter with hadron matter; (2) optimized state-of-the-art solvers for lattice QCD (Quantum Chromodynamics) with Bayesian optimization techniques; (3) studied scaling laws of distributed DL approaches on Polaris; (4) developed and optimized benchmark suite for Intel extension of PyTorch for enabling DL simulations on Aurora; (5) developed infrastructure for machine learning model operationalization management (MLOps) at LCF; (6) published two articles in *Phys. Rev. D*, gave multiple invited talks including the plenary talk at a SIAM workshop, and gave a talk and tutorial at ALCF's 2022 Computational Performance Workshop.

Nick Frontiere (Ph.D., physics, University of Chicago). **Hired:** April 2019. **Terminated:** November 2022. Transferred to another Argonne division in a staff role in December 2022. **Research area:** Numerical cosmology. **Projects:** ECP ExaSky project (PI Salman Habib) and lead PI for 2022 LDRD (Laboratory-Directed Research and Development) project. **Accomplishments:** (1) Development of a GPU-optimized Hydrodynamic Solver in the simulation code HACC; (2) co-design and preparation work for the deployments of Aurora and Frontier: these efforts involved porting the HACC CUDA solvers to HIP and SYCL/DPC++ languages and performing vendor specific tuning; (3) developed several scientific capabilities to improve the fidelity of cosmological simulations, including development and calibration of standard sub-resolution ("subgrid") model prescriptions into CRK-HACC. The full new "CRK-HACC" framework was described in a recently accepted submission to *ApJ* earlier this year. This work also included extensive validation measurements of idealized hydro and gravitational tests to prove efficacy, with several detailed comparisons with the AMR code Nyx discussed in submission to *Monthly Notices of the Royal Astronomical Society (MNRAS)* preprint);

(4) published a journal article on a high-resolution cosmology simulation at the gigaparsec scale in *ApJS* **259**.

Raymundo Hernandez-Esparza (Ph.D., chemical sciences, Metropolitan Autonomous University, Mexico City, Mexico). **Hired:** September 2022. **Research area:** Computational chemistry. **Project:** Chemical catalysis at exascale. **Scientific goal:** To measure the performance of density-functional theory (DFT) code Quantum Espresso in systems of interest for heterogeneous catalysis. **Accomplishments:** Investigated the performance of Quantum Espresso for simulating gas-phase heterogeneous catalysis using both GPU-aware MPI and multi-process service (MPS) interface on Polaris and found the execution times were reduced by up to 25 percent. These results and other developments on workflow tools were presented at the 2022 annual meeting of the Exascale Catalytic Chemistry (ECC) project.

Chunyong Jung (Ph.D., atmospheric sciences, North Carolina State University). **Hired:** December 2021. **Research area:** Regional climate modeling, atmospheric dynamics, high-impact weather systems (e.g., tropical/extratropical cyclone). **Current projects:** ALCC, LDRD (Towards Neighborhood Scale Climate Simulations using AI and Accelerated GPUs). **Scientific goal:** To better understand climate change effects on high-impact weather events, such as heavy precipitation systems and tropical cyclones and their transition to extratropical cyclones through high-resolution idealized modeling. **Accomplishments:** Convection-permitting historical and future regional climate dataset that covers entire North America was built, leading to a winning INCITE proposal. The dataset was extensively analyzed and subsequently presented at various academic conferences.

Geng Liu (Ph.D., mechanical engineering, City College of New York). **Hired:** October 2021. **Research area:** Parallel computing. **Current projects:** Aurora ESP project *Extreme-Scale In-Situ Visualization and Analysis of Fluid-Structure-Interaction Simulations* (PI: Amanda Randles). **Scientific goal:** advancing the use of data science to drive in situ visualization of extreme-scale fluid-structure-interaction simulations. **Accomplishments:** (1) Presented conference poster “Case Study for Performance Portability of Lattice Boltzmann Kernels” at SC22; (2) working on HARVEY, the code used in the lattice Boltzmann project, as well as its proxy application. In the original HARVEY code, the programming model for GPU parallel computing is CUDA. The team successfully translated the simulation part of HARVEY to the Kokkos version and developed the Kokkos and SYCL versions of the proxy application based on the original CUDA version.

Nathan Nichols (Ph.D., material sciences, University of Vermont). **Hired:** September 2021. **Research area:** Quantum simulation of low dimensional systems, QMC (Quantum Monte Carlo) algorithmic development, analytic continuation, ML for science. **Project:** ATLAS Aurora ESP program. **Scientific goal:** To prepare HEP event-generation code used by the ATLAS Experiment for the exascale era of computing. **Accomplishments:** (1) Wrote a SYCL port of MadGraph4GPU (MG) event generation code and assisted in updating the Kokkos port and led an effort to integrate the SYCL and Kokkos C++ code generation plugins and libraries into the full application. The SYCL version achieves improved performance over the native CUDA implementation on NVIDIA devices while maintaining portability to devices from other vendors. Results were presented at the 41st International Conference on High Energy Physics (ICHEP

2022) and the 21st International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT 2022); (2) set up CI (Continuous Integration) scripts that track performance of MG for several JLSE (Joint Laboratory for System Evaluation) testbeds, collecting data on AMD, NVIDIA, and Intel experimental hardware for five HEP processes of increasing computational complexity. The CI records figures-of-merit for CUDA, OpenMP, Kokkos, SYCL, and Alpaka builds of MG; (3) released open-source software to perform the analytic continuation of imaginary time correlation functions using differential evolution which is already seeing use across the QMC community; (4) research into low dimensional exotic quantum phenomena has produced an experimental realization of one-dimensional helium that was presented at the APS March Meeting and published in *Nature Communications*. Work on a strain-induced superfluid-insulator transition for atoms adsorbed on graphene was also presented at the APS March Meeting; (5) taught Python at the ACT-SO Summer Coding Series, a 6-week course to teach coding to middle and high school students from underserved communities.

Nwamaka (Amaka) Okafor (Ph.D., electrical and electronic engineering, University College Dublin). **Hired:** November 2022. **Research area:** Applying AI and Internet of Things (IoT) to environmental monitoring. **Current projects:** System log analysis, including sensor data, hardware errors, and job logs.

Siddhisanket (Sid) Raskar (Ph.D., electrical and computer engineering, University of Delaware). **Hired:** June 2021. **Research areas:** dataflow models and architectures, HPC, ML, AI accelerators, performance modeling. **Scientific goal:** Evaluating the efficacy of AI architectures for scientific machine learning and the design of next-generation AI architectures for science. **Accomplishments:** (1) Participated in the MLPerf HPC submission from Argonne; (2) evaluated AI Testbeds (performance modeling, benchmarking); optimized DNN on heterogeneous architectures; (3) evaluated GNN (Graph Neural Network) operations and models on AI Testbeds; (4) served as co-PI of a LDRD Expedition to evaluate ALCF's Graphcore accelerator; (5) produced five publications (two published and three under review) and gave invited talks at ALCF training events, a tutorial at SC22, and a lighting talk at the 2022 LLVM Developers' Meeting.

Mathi Thavappiragasam (dual Ph.D. in electrical and computer engineering (ECE) and computational mathematics, science and engineering (CMSE), Michigan State University). **Hired:** June 2022. **Research area:** Heterogeneous frameworks and low-level optimization. **Current projects:** (1) ExaStar (porting and optimizing ExaStar applications on Aurora); (2) MiniMDock (porting this particle-grid-based molecular docking application); (3) high-level heterogeneous programming model research (studying performance portability capabilities of OpenMP and SYCL using different applications and data transfer strategies between host and device to challenge with computing speed of the accelerators). **Scientific goal:** Enhancing the performance and portability of application software using novel architecture features. Studies to form specifications for the features of future architectures. **Accomplishments:** (1) Building ExaStar-Thornado on Intel GPU systems with support of the Intel-Argonne COE, Intel exascale performance codesign team and Intel compiler teams, and improved the oneAPI compiler features in the process; (2) discovered untapped performance gains in compiler configurations and runtime environmental variables and improved the performance of the dominating kernels significantly (~10X) and achieved 50 percent speedup on Intel PVC over NVIDIA A100;

(3) successfully ported MiniMDock on Intel systems using OpenMP offloading features and tuned the performance based on Intel oneAPI compiler features; (4) gave a contributed talk at HiPar-22 at SC22 that motivated the team to elaborate the study and submit a journal article; (5) submitted a proposal for a SULI internship for summer 2023 to work on a benchmark analysis over different systems/compilers using OpenMP; (6) studying the reengineering of numerical schemes for initial boundary value problems to derive more compute intensive solutions that focus on steering away from costly sparse matrix operations when possible.

Umesh Unnikrishnan (Ph.D., aerospace engineering, Georgia Institute of Technology). **Hired:** November 2021. **Research area:** HPC, performance engineering, software portability, CFD. **Scientific goal:** To develop a scientific application targeting CFD applications of high-speed compressible flows that is performant and portable to Aurora and other upcoming exascale architectures. **Current projects:** (1) Development of the CNS-libParanumal application for high-speed applications – investigation of shock capturing schemes and reacting flow capabilities; (2) performance benchmarking of ECP-CEED, libParanumal software library and associated kernels on Aurora developmental hardware; (3) investigation and implementation of optimization techniques for compute-intensive kernels in the CNS-libParanumal application. **Accomplishments:** (1) Gave three conference talks; completed code improvements to libParanumal code base to run with the Intel oneAPI SDK on Aurora hardware; (2) developed mini-apps for benchmarking and testing optimized versions of key kernels of the CNS application; (3) developed modified versions of the most expensive kernels that have shown a speedup of 1.3-1.7x over the original kernels; (4) developed mini-apps for developing and testing new science features such as Burgers' equation solver and advection-diffusion-reaction solver.

Madhurima Vardhan (Ph.D., biomedical engineering, Duke University). **Hired:** September 2022. **Award:** 2022–2023 Margaret Butler Fellow. **Research area:** Biomedicine. **Current projects:** (1) Developing scalable AI algorithms for predicting risk of cardiovascular diseases and Long Covid trained using synthetic datasets that resemble real patient data; (2) assessing use of genetic data in predicting major adverse cardiac events. **Accomplishments:** (1) Helped set up cross collaborations with clinical and biomedical researchers at the University of Illinois at Chicago (UIC), U.S. Department of Veterans Affairs, and the University of Chicago (UC) to work on predicting the 10-year cardiovascular disease risk and Long Covid; (2) set up collaboration with the Clinical Practice Research Datalink at the UK Biobank to access high-fidelity synthetic datasets for cardiovascular diseases and Long Covid to develop AI models using ALCF resources to capture complex data distributions; (3) participated in the NIH N3C Rapid Acceleration of Diagnostics (RADx) challenge; (4) holds joint appointments at UC and UIC to facilitate collaborations in the intersection of HPC/AI and biomedicine; (5) completed all training required to perform Human Subject Research.

Kaushik Velusamy (Ph.D., computer science, University of Maryland, Baltimore). **Hired:** January 2022. **Research area:** Data management, storage, and I/O. **Project:** ECP project ExaIO to develop efficient I/O libraries for exascale computing. **Scientific goal:** Developing a topology-aware custom collective I/O (CCIO) in Hierarchical Data Format version 5 (HDF5), one of the most popular scientific I/O libraries, to improve the I/O performance of various ECP and ALCF applications. **Accomplishments:** (1) performed evaluation studies on Theta, Polaris, and Summit using a suite of parallel I/O benchmarks called *h5bench*, which represents the I/O patterns commonly used in HDF5 applications on HPC systems. The initial performance evaluation shows ~5x performance improvement compared to baseline HDF5; (2) preparing a paper on the design and implementation of CCIO, and performance evaluation results on Theta, Polaris, and Summit; (3) presented a talk and hands-on session at ATPESC and other ALCF-hosted training workshops; and (4) volunteered for various outreach efforts aimed at high school students (Argonne summer research training, ACT-SO 2022) and service to Argonne's Postdoctoral Society.

Azton Wells (Ph.D., physics, University of California, San Diego). **Hired:** September 2022. **Research area:** Cosmology, data analysis, deep learning. **Current projects:** (1) HACC cosmology code frontend development; (2) exploring DL methods to learn data-driven models of the UV background of the universe using simulations of the Lyman alpha forest. **Accomplishments:** (1) Incorporated HACC frontends into the astrophysical analysis API, YT. This change will make a host of simple, Python-based analysis tools available to scientists who are not part of the core HACC development team, as well as expedited visualization capabilities for validation of our methods and tuning simulations. Also added frontend integration for halo, galaxy, and core catalogs; (2) used DL methods to perturb the Lyman-alpha forest statistics in HACC simulations so that the team may reproduce observable Lyman-alpha statistics, such as the flux probability distribution. The DL model creates this prediction by learning the UV flux given the simulated densities and temperatures; (3) finalized dissertation into a paper submission to *ApJ*.

Zhen Xie (Ph.D., Computer Science and Technology, The Institute of Computing Technology of the Chinese Academy of Sciences, Beijing, China). **Hired:** August 2021. **Research area:** Mixed precision optimization for next-generation AI accelerator systems. **Current projects:** (1) Test different DNN microbenchmark and communication functions on each platform in a first-of-its-kind evaluation of diverse DNN workloads and collective communication on AI accelerators, with opportunities to integrate these novel AI accelerators in supercomputing systems; (2) build genome-scale language models (GenSLMs) that can learn the evolutionary landscape of SARS-CoV-2 genomes and demonstrate the scaling of GenSLMs on both GPU-based supercomputers and AI-hardware accelerators, achieving over 1.54 zettaflops in training runs; (3) develop a high-performance DNN training engine, which ensures sufficient model accuracy as single-precision training while maximizing the performance benefits of half-precision format, BFloat16; (4) test and optimize the opencatalyst-benchmark on the Polaris system for the MLPerf training benchmark project suite to measure how fast systems can train models to a target quality metric; (5) mentor and advise intern students on an LDRD project, a scalability study of AI-based surrogate for ptychographic image reconstruction on Graphcore and Habana systems. **Accomplishments:** (1) Produced two publications (one published and one in review) and gave talks at the 13th IEEE Performance Modeling, Benchmarking and Simulation (PMBS) at SC22 and at the 27th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming (PPoPP); (2) won a Gordon Bell Special Prize for HPC-Based COVID-19 Research.

Section 5. Risk Management

(a) Does the Facility demonstrate effective risk management practices?

ALCF Response

The overview of the risk management process that ALCF follows (and laid out in Section 5.1) clearly demonstrates that ALCF successfully managed both its project risks and operational risks in 2022. 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. ALCF currently has 37 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 2022 are listed in Section 5.2, along with the details of the risks in Table 5.1. Events stemming from those risks and the mitigations for them are described in greater detail in Section 5.3. Section 5.6 and Table 5.1 provide details on the major risks that will be tracked in 2023.

5.1 Risk Management Process Overview

ALCF uses documented risk management processes, first implemented in June 2006, and outlined in its RMP, for both operational and project risk management. The RMP 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.

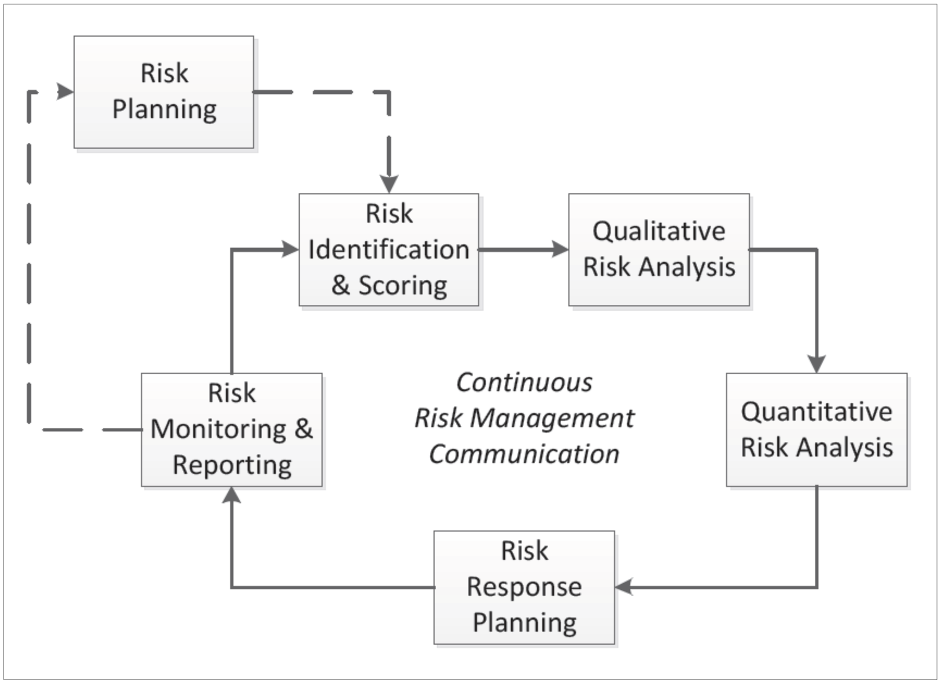

Figure 5.1 Illustration of the Iterative Risk Management Process from ALCF's Risk Management Plan

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

1. Plan, implement, and revise the RMP.
2. Identify threats and opportunities to 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 and mitigate the risks.
4. Monitor risks, mitigation plans, and management reserve and contingency until the risks are retired or the project is closed.

A key part 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; risk forms and the risk register are formatted using Excel. Risk owners continually monitor the risks they own and submit monthly reports on all risks through the ALCF online risk reporting system.

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. The board serves 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 whether to add proposed risks 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 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, such as a change in the technology landscape.
  - Works with the risk owner to modify the risk statement if needed to risk mitigation strategies, risk triggers, or risk scope.
  - Decide if a risk ownership should change.
- Reviews and identifies any risks to retire.
- Reviews the risks encountered in the past 18 months to discuss potential actions.
- Discusses risks encountered at other facilities and how they might apply to ALCF.

5.1.2 Risk Management in Day-to-Day Operations

ALCF currently has 37 open risks in the facility operations risk register and uses post-mitigated risk scoring to score the risks. The 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), so risks are distributed throughout the division. Risks are owned by group leads or managers. In ALCF operations, subject matter experts estimate risk mitigation costs for management consideration. In addition to formal Risk Review Board meetings, ALCF has many formal and informal individual risk meetings as needed. The Risk Review Board did not meet in CY 2022.

Risks are identified and evaluated, and mitigation actions are developed for all changes that occur at ALCF—from installing a new piece of hardware to changing the scheduling policy to upgrading software. For example, new risks are created anytime a resource goes from project to steady-state, or when a resource is decommissioning. If the risks are short term or minor, they are not added to the registry. New significant risks identified during the individual meetings are added to the registry and reviewed at the next Risk Review Board meeting. Any change made to an existing risk—whether recharacterizing it, changing the factors affecting it, or retiring it entirely—requires a formal meeting, which can be proposed by any risk owners.

Other tools besides the risk register are used for managing risks in day-to-day operational risk. 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 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 identify 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 guide the work.

In addition to machine operations, risk management is used in such diverse ways as evaluating and managing INCITE and ALCC proposal risks (e.g., 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 (e.g., 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 ALCF's next supercomputer, known as ALCF-3, continued in 2022. Risk Register managers continue to maintain a project risk register and track a set of detailed risks. Risk mitigation costs are developed using a bottom-up cost analysis, then are input to the commercial project risk analysis tool Oracle Primavera Risk Analysis (OPRA) to manage the contingency pool. These risks are not included in the risk numbers covered in this document and are not discussed further.

5.2 Major Risks Tracked

The ALCF operated both Theta and Polaris during 2022 and planned the growth of both the staff and the budget to bring the facility to full strength. As such, ALCF continues to monitor a number of major risks for the facility. No risks were retired in 2022.

Four major operations risks were tracked for 2022, two with a risk rating of High, one Moderate, and one Low. None of these were encountered, and all of them were managed. They are described in Table 5.1. All risk ratings shown are post-mitigation. The risks are color-coded as follows:

- Red risks are Moderate or High.
- Orange risks are Low.

Table 5.1 Major Risks Tracked for CY 2022

ID	Title	Encountered	Rating	Notes
1059	Funding/Budget Shortfalls	No	High	ALCF regularly worked with the program office to plan a budget for handling the impact of a Continuing Resolution in FY 2023, new hires, and changes in the laboratory indirect expense rate. This risk remains a major concern as the facility moves forward with Theta, Polaris, and ALCF-3 in CY 2023.
25	Staff Recruitment Challenges	No	High	The ALCF added six new staff members in CY 2022: five of those were external new hires and one is an employee reclassified from postdoc to staff. 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	No	Moderate	ALCF lost three staff members during CY 2022, all of whom terminated employment at Argonne. Budget concerns at Argonne and the growth in high-paying industry jobs for system administrators and programmers with HPC expertise make staff retention in future years a continuing concern.
1091	Injury to Workers/Overall Safety of the Division	No	Low	ALCF had increased data center activity in CY 2022, but zero incidents given the safety culture of ALCF.

5.3 Risks Encountered in the Review Year and Their Mitigations

ALCF encountered five risks in CY 2022. The risk owners are identified below, along with an assessment of the risk's probability and impacts, a brief description of what transpired, and how the risk was ultimately managed. All five risks have a residual rating of Very Low.

5.3.1 Interruptions to Facility that Provides Cooling

0030: Interruptions to Facility that Provides Cooling	
Risk Owner	Director of Operations
Probability	Very Low
Impact	Cost: Very Low; Technical Scope: Very Low
Risk Rating	Very Low
Primary Management Strategies	Increase redundancy. Site cooling piping is now interconnected to enable multiple chiller plants to back up each other. Implement emergency shutdown software to limit impact of increased heat on hardware and to reduce heat generation. Initiate root cause analysis on any failures or degradation and remediate accordingly.
Triggers	Rising temperatures on equipment in the Data Center; planned maintenance; monitoring notification of outages or rising temperatures.

Description

On August 1, 2022, the TCS Building (TCSB) management company JLL received an alarm at 12:30 p.m. and discovered that Chiller 1 had shut down. Following an unsuccessful attempt to restart the chiller at 2:00 p.m., JLL called in a third-party mechanic. At 3:45 p.m., Chiller 2 shut down due to overload, despite ALCF attempts to reduce the heat load.

Evaluation

At 4:15 p.m., the mechanic determined that the reservoir outlet hose on Chiller 1 had come off, draining the reservoir, and causing the shutdown. Because the setup is not N+1, Chiller 2 overloaded and shut down. Chiller 1 was successfully restarted at 4:50 p.m. Chiller 2 was restarted at 5:30 p.m. but shut down again after 15 minutes, due to overload. The mechanic adjusted the controls and was able to restart it at 6:30 p.m.

Management

Argonne and the TCSB Trust evaluated options to improve cooling capacity and redundancy.

5.3.2 Interruptions to the ALCF Network

0032: Interruptions to the Facility Network	
Risk Owner	Systems and Networking Team Lead
Probability	Very Low
Impact	Cost: Very Low; Technical Scope: Very Low
Risk Rating	Very Low
Primary Management Strategies	Argonne's Business and Information Services (BIS) Division will coordinate any planned maintenance or modifications. Initiate root cause analysis on any failures, or degradation, remediate accordingly.
Triggers	Planned maintenance; Failure of all the redundant paths.

Description

On June 21, 2022, Argonne's campus border routing infrastructure suffered an issue that disrupted networking. This issue propagated to the Argonne campus router serving ALCF's data center, resulting in intermittent connectivity issues for ALCF resources.

Evaluation

Argonne updated the campus border router code to address the issue on the same day. ALCF installed a configuration to guard against this issue with the BIS campus router.

Management

ALCF asked Argonne's campus networking group to diversify the devices serving the ALCF network.

5.3.3 Signing Delays of Some INCITE/ALCC User Agreements

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

Description

The 2022 INCITE project “TropicalMeteorology - New Window into Tropical Meteorology with Global 1 km Atmosphere-Ocean Simulations” was granted a 50K node-hours allocation on Polaris, to start on August 9, 2022. A signed Master User Agreement (MUA) between the project team’s institution, European Centre for Medium-Range Weather Forecasts (ECMWF), and Argonne needed to be in place before the project team members were allowed access to the system. The ECMWF team requested updates to the Nonproprietary (NP) User Agreement which resulted in a delay for Argonne lawyers to consider the change.

Evaluation

ECMWF requested additional language in Article IX of the NP User Agreement; specifically, clarification on what constitutes technical data and whether their proprietary code can be excluded from it. The ECMWF team had previously requested a Proprietary User Agreement but upon reviewing that agreement’s terms and conditions decided that an NP User Agreement with updated language was more appropriate. ALCF’s user experience team lead, industry outreach manager, and director of science reviewed the matter and reached out to Argonne Legal for further guidance.

Management

Argonne Legal stated that the NP User Agreement template is common across all projects at all Argonne user facilities and tailoring the language per project and/or per facility would lead to a significant overhead cost and delay. Argonne Legal recommended using a Nondisclosure Agreement (NDA) in conjunction with a standard NP User Agreement and incorporate information regarding the proprietary code and a few other additional points. This recommendation was conveyed to the ECMWF team and, after several iterations with their legal team and the Argonne legal team, this was the path used to get to a signed user Agreement in place between the two organizations.

5.3.4 System Stability Issues

1078: System Stability Issues	
Risk Owner	Systems and Networking Team Lead
Probability	Very Low
Impact	Cost: Very Low; Technical Scope: Low
Risk Rating	Very Low
Primary Management Strategies	Conduct regression testing to monitor system health. Monitor and analyze system failure data. Activate swat team including vendor and ALCF Operations staff to do root cause analysis and develop workaround/resolve problem.
Triggers	User complaints. System failure rate increasing.

Description

On February 4, 2022, problems with interconnect links caused instability on Theta resulting in an unplanned outage.

Evaluation

Normally these links are redundant, but enough problem links were present that an unrecoverable failure occurred.

Management

Newly hired staff systems administrators missed running scripts to detect troubled links. This was corrected through appropriate staff training.

5.3.5 Diagnostic Suite and Utilities Fail to Detect Hardware Problems

1085: Diagnostic Suite and Utilities Fail to Detect Hardware Problems	
Risk Owner	Systems and Networking Team Lead
Probability	Very Low
Impact	Cost: Very Low; Technical Scope: Very Low
Risk Rating	Very Low
Primary Management Strategies	Track and monitor job and hardware failures, correlate. Work with vendor to resolve issues and improve diagnostic suite. Activate swat team to work with the vendor and user to determine root cause and to develop workaround/resolve the problem.
Triggers	Job failures increase and hardware failures do not. User complaints.

Description

On February 4, 2022, problems with interconnect links caused instability on Theta resulting in an unplanned outage.

Evaluation

An automated diagnostic utility was missing a search term meant to detect problem interconnect links on Theta. This search term is normally present but had been removed inadvertently.

Management

The search term was re-added to the Simple Event Correlator (SEC) that analyzes logs and notifies systems administrators of issues on Theta. The SEC configuration was also reviewed for correctness.

5.4 Retired Risks

No risks were retired in CY 2022.

5.5 New and Recharacterized Risks

There are no new risks and no recharacterized risks to report in 2022; however, the bug incident resulting from a Lustre upgrade to ALCF e1000 storage systems in September and fixed without any found instances of data corruption in December, may warrant further monitoring. The risk management team is discussing the creation of a new risk in response to this.

5.6 Top Operating Risks Monitored Closely for the Next Year

Table 5.2 lists the current top operating risks that will be closely monitored in CY 2023, 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 Top Operating Risks Monitored for CY 2023

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.
0990	Electricity cost could increase beyond the budget	Low	Contract negotiations by laboratory management to lock in long-term contract. Monitor electricity usage data. Evaluate methods to improve equipment efficiency and lower electricity usage. Consider activating the power throttling feature on the hardware to reduce electricity usage, at the expense of slowing down the processors and lowering machine performance. Include escalation.

Conclusion

ALCF uses a proven risk management strategy that is documented in the RMP. Risks are reviewed and updated to reflect the dynamic nature of risk management, as well as to document new lessons learned and best practices captured from other facilities. Risk management is part of ALCF's culture and used by all staff, from senior management to summer students. A formal risk assessment is performed for every major activity within ALCF, with informal assessments used for smaller activities. Risks are monitored and tracked using a secure and shared cloud-based storage system. Many other tools are used to manage risks at ALCF, particularly in the area of safety. ALCF's effective risk management plan has contributed to the successful management of all significant risks encountered in CY 2022.

Section 6. Environment, Safety, and Health

(a) Has the Facility demonstrated effective Environment, Safety, and Health (ES&H) practices to benefit staff, users, the public, and the environment?

ALCF Response

ALCF continues to be a leader in its commitment to health and safety. ALCF is guided by Argonne's Worker Safety and Health Program, which established the framework for using the Integrated Safety Management (ISM) system to comply with 10 CFR 851. ALCF has developed and implemented directorate- and division-level policies and procedures consistent with Laboratory directives to support safe and effective work.

6.1 Work Planning

ALCF work activities are planned, reviewed, and documented in accordance with its divisional work planning and control (WPC) manual. The WPC manual incorporates the guidance from DOE-HDBK-1211-2014: *Activity-Level Work Planning and Control Implementation* and details how ALCF executes ISM principles in planning and execution of work.

ALCF work activities are reviewed and documented using two methods: Work Control Documents (WCDs) and Job Safety Analysis (JSA). ALCF has a main WCD called *Operations Skill of the Worker Standard Tasks* that is used to describe what, how, and where the most common low-rigor operations tasks in the data center occur. ALCF also uses specific task-based WCDs to plan for work that has a limited scope, is a unique activity, or requires a specific skill set. ALCF currently has a specific WCD to cover replacement of rectifiers that requires a Qualified Electrical Worker (QEW) to perform the work. ALCF uses a graded approach for approving and authorizing WCDs: More complex or hazardous work brings higher levels of involvement from laboratory management or subject matter experts (SMEs) into the approval and authorization process. WCDs are reviewed and updated annually to ensure that they capture all the hazards and controls workers might encounter in the data center. New workers are added to the WCD as authorized workers once they have completed appropriate training and demonstrated proficiency. Feedback from workers and from the yearly reviews are incorporated into the revised WCD to complete the ISM cycle.

The second method ALCF uses to capture and review work activities is a JSA. In CY 2022, the division developed 16 JSAs to cover non-Argonne contractors and vendors who perform maintenance, troubleshooting, and repair work in the data center. Each JSA describes how ALCF, in collaboration with the contractors, plans and identifies hazards and controls for the work activities that the contractor will perform. ALCF is responsible for all contractors performing work, so works closely to ensure that all work is completed within the scope limits and follows the hazards mitigations outlined in the JSA. Like the WCD, the JSA uses a graded approach, and senior management and SMEs are engaged for higher-risk or more complex work.

6.2 COVID-19 Site Reentry

In CY 2022, Argonne moved from limited operations with maximum telework and transitioned to normal operations with optional telework. ALCF's reentry plan maintained flexibility, authorizing its employees to continue to telework all or a percentage of their time post-reentry, with their supervisor's approval. ALCF resumed doing in-person ergonomic evaluations, completing four such evaluations in CY 2022, while also continuing to provide virtual ergonomic evaluations.

ALCF acknowledged and addressed an unintended risk to workers with fewer people working onsite. In CY 2022, data center emergency procedures were updated to better protect staff working in a reduced-occupancy environment and to streamline the JSA used for planning contractor work. Argonne's reentry plan included increased or reduced controls based on COVID-19 community levels. The JSA updates included recommendations to determine the impact of COVID-19 on the community and to act, with the flexibility to change controls for COVID-19 without a complete revision of the JSA.

6.3 Continued Operations

ALCF uses the plan-do-check-act (PDSA) cycle as a model to incorporate ISM functions, wherein management assessments are performed as the feedback portion of the ISM principles. These assessments are planned out at a directorate level based on risk, adjusted if needed, and repeated again for continuous improvement. In CY 2022, ALCF updated its tour guide training and guidance based on findings from the Noise in the Data Center Assessment.

ALCF continues to conduct biannual health and safety inspections led by its Division Director and Environmental, Safety and Health (ESH) Coordinator. These inspections serve as an additional opportunity to seek feedback from staff. Feedback is also captured at the Operations team's biweekly all-hands meetings, which begin with a safety discussion and are regularly attended by the ESH Coordinator and Data Center Manager to ensure that safety and work plans are communicated.

6.4 Aurora Deliveries

The CY 2022 delivery of the Aurora computer and non-compute racks occurred without incident. The delivery period began November 9, 2021, and ended on March 8, 2022, and consisted of the delivery and placement of 166 compute racks, 76 non-compute racks, 57 vCDUs (virtual Channel Data Units), and 24 PDUs (power distribution units) in the data center. Management of all system component deliveries and rack placement was planned, executed, and coordinated by ALCF staff, the Project Management Organization (PMO), TCS building management, and the subcontractors.

- A JSA was developed by the subcontractors and the Aurora Floor Manager using the initial Aurora project hazard analysis and input from laboratory SMEs.
- A schedule and move plan were developed by the Aurora Project Manager and Coordinator in conjunction with the TCS building's owner and construction manager, Argonne SMEs, and subcontractors.

- All subcontractors were trained using the JSA, on building orientation, the ESH Plan for Working in the Building 240 Data Center, and on the current COVID-19 controls prior to being authorized to work. Each new subcontractor received this training.
- Daily pre- and post-job meetings were conducted with the subcontractors, Aurora Project Manager, Aurora Floor Manager, all TCS building personnel involved, ALCF's ESH Coordinator, and ALCF support staff.
- Three safety pauses occurred during the delivery period; two of them occurred in CY 2022. They address the following concerns:
    - On December 8, 2021, the PDU crate ramps failed. The subcontractor created reinforced ramps to be used moving forward.
    - On January 12, 2022, crate doors were leaned against a wall. Argonne created space to store doors safely moving forward.
    - On January 12, 2022, a subcontractor incorrectly used a ladder. The subcontractor and Argonne staff practiced safe ladder use at the morning safety brief.
    - On January 18, 2022, a non-compute rack crate hit duct work in the loading dock. TCS building personnel taped out a floor path to avoid area duct work.

6.5 Incident Reports

In CY 2022, ALCF experienced zero Days Away Restricted or Transferred (DART), zero first aid cases and one Total Recordable Case (TRC).

Table 6.1 ALCF Injury and Illness Cases FY 2018 thru FY 2022

	DART	First Aid	TRC
2018	0	1	0
2019	0	1	0
2020	0	0	0
2021	0	0	0
2022	0	0	1

In CY 2022, an ALCF staff member fell in a stairwell and suffered a broken wrist. A subsequent incident investigation found the stairwell to be free of hazards or unexpected conditions. The employee was descending the stairs while holding safety glasses in one hand and a cellphone in the other and missed the last two steps and fell forward. Medical personnel in Argonne's Health and Wellness Clinic examined the employee and referred her to a medical facility offsite where X-rays confirmed a wrist fracture. The employee was given restrictions that did not inhibit her ability to perform her job. The injury was determined to be a DOE Occurrence Reporting and Processing Service (ORPS) incident and recorded, but resulted in no lost days. A corrective action was developed with the assistance of the Argonne Incident Investigation Team that included stairwell safety awareness signs that were posted across Argonne.

This page intentionally left blank.

Section 7. Security

(a) Has the Facility demonstrated effective cyber security practices?

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

(c) Does the Facility have effective processes for compliance with applicable national security policies related to Export Controls and foreign visitor access?

ALCF Response

ALCF works to continually improve its cyber security practices by developing and maintaining relationships between facility personnel and its internal security personnel, and between ALCF and the Argonne Cyber Security Program Office (CSPO). ALCF is committed to improving relationships with sister laboratories, auditing new software deployments, staying up to date on emerging cyber threats and attacks, and addressing future facility needs as they relate to security controls. This resulted in ALCF having zero security incidents this year and allowed ALCF to reduce security risks. ALCF has a valid Authority to Operate (ATO) and an effective process for compliance with export controls and foreign visitor access.

7.1 Continual Improvement in Cyber Security Practices

ALCF maintains a strong cyber security posture and is vigilant in monitoring active and evolving threats, vulnerabilities, and potential exploitation. ALCF's security team partners with OLCF, NERSC, and other facilities to quickly assess the risks of security incidents unique to HPC facilities. Working together, the facilities can quickly identify, coordinate, and verify mitigations for new vulnerabilities on HPC systems.

In 2022, security team members from ALCF, OLCF, NERSC, and ESnet formalized the Secure ASCR Facilities (SECAF) working group and established a steering committee made up of one security representative from each organization. With the support of the facility directors and DOE, SECAF members adopted a charter and began actively collaborating in spring 2022. SECAF's purpose is to build community and provide opportunities for cross-laboratory collaboration to improve site security. SECAF's cross-laboratory communications and quarterly full-team syncs will be supported by DOE Mattermost server.

Additionally, ALCF's security staff worked to improve workflows for users, especially by identifying and eliminating potential procedural barriers. In CY 2022, ALCF continued to examine user workflow improvements and new tools, both locally and at the laboratory level, to increase cyber security controls and to reduce the manual efforts involved with regular tasks. ALCF's security staff and CSPO staff have a good working relationship and regularly discussed future initiatives.

In CY 2022, there were zero cyber security incidents on ALCF systems. ALCF's cyber security personnel take a proactive approach to problem management. Examples include efforts to:

1. Conduct privileged access reviews across the environment to help ensure that everyone has the appropriate level of access.
2. Conduct reviews of how and where data are stored to help ensure that data are accessible only to those with proper authorization.
3. Educate users and staff about how to prevent password exposure.
4. Educate developers on secure coding best practices via internal discussions/reviews and external training courses.
5. Integrate security auditing into developer workflows to identify security issues early in the development life cycle.
6. Keep the National Institute of Standards and Technology (NIST) certification package up to date, including NIST 800-53, 800-34, 800-30, and 800-18 compliance documents.
7. Archive and delete obsolete data.
8. Set password rotation policies for ALCF systems and verify compliance.
9. Monitor new vulnerabilities to ALCF systems.
10. Conduct penetration testing of both internal- and external-facing web applications and recommend security improvements.

Additionally, ALCF takes a layered approach to security in areas that fall within Argonne's CSPO-managed cyber security domain. The laboratory makes the following security services available to ALCF:

1. Cross-laboratory data sharing that CSPO integrates into automatic network blocking rules to keep systems secure.
2. Log analysis capabilities that allow CSPO to block IP addresses or bad actors in near real-time. This provides ALCF additional protection from attacks on Argonne networks as they are detected.
3. CSPO provides Cloudflare, a powerful web application firewall (WAF), which protects ALCF's public sites from known attacks and ensures encryption requirements are met.
4. CSPO provides a Tenable security center (Tenable.sc) instance that ALCF can tie Nessus scanners into and allows leveraging CSPO's public and internal network scanners. This provides a holistic view of network vulnerabilities in ALCF's environment.
5. A laboratory-managed network border firewall protects ALCF applications from the public Internet and networks within Argonne.
6. CSPO helps to manage publicly visible vulnerabilities by automatically alerting ALCF when new vulnerabilities are detected and provides guidance for patching the system or removing public conduits.
7. CSPO helps with cyber security policy by managing the Cyber Security Program Plan and allows ALCF to make documented alternative security implementations to suit ALCF's needs.
8. Access to NetFlow data to help troubleshoot and investigate network-related issues across the laboratory's networks.
9. Access to cybersecurity tools such as Axonius for asset management and maintaining an inventory of installed software.

Some ALCF proactive measures revealed security vulnerabilities that were promptly addressed and fixed, usually within days of their discovery. Immediately after detection, ALCF staff investigated all relevant logs to determine whether the security vulnerability had been exploited. In 2022, none of the issues that were investigated were found to have been exploited. Examples of the security issues that were detected, and their ensuing mitigations, are as follows:

1. **Issue:** The CSPO identified new security issues with public-facing ALCF services.
   Mitigation: ALCF staff reviewed the CSPO information and worked to address the issues in a timely manner.
2. **Issue:** Passwords in some applications were found to be stored insecurely.
   Mitigation: Evaluation of the application data integrity showed no unauthorized access. Passwords were changed as a precaution against potential exploit.
3. **Issue:** CSPO received a report from the U.S. Department of Homeland Security (DHS) via the Integrated Joint Cybersecurity Coordination Center (iJC3) that insecure SSL ciphers were configured for the Globus data transfer nodes at ALCF.
   Mitigation: ALCF staff worked with Globus engineers to find the root cause and release a new fix to bring ciphers in line with DHS expectations. ALCF also informed other sites through SECAF to help ensure they were also able to address the issues.

ALCF will continue to proactively investigate security issues and to monitor and respond to all vulnerabilities. Plans for improving the security of ALCF resources include:

1. Retiring obsolete services and data.
2. Verifying that strong encryption is used everywhere in the ALCF environment and that plain text protocols are not used for production systems.
3. Improving real-time log analysis techniques.
4. Increasing the visibility into required security updates on systems.

CSPO conducts an annual internal audit called a Division Site Assist Visit (DSAV) to assess divisional compliance with NIST-800-53 controls. Each DSAV covers roughly one-third of the controls. In CY 2022, CSPO's DSAV with ALCF concerned 40 NIST controls. Of these, CSPO found that ALCF has 31 controls that are fully implemented within policy, has partially implemented 8 others that have room for growth, and has one control that needs attention. ALCF will review the 9 controls that were not recognized as fully implemented with policy and assess ways that we can improve over the next year. The opportunities for improvement identified by CSPO are as follows:

1. Develop automated means by which flaw-fix updates are identified, evaluated, and prioritized for applicability within the environment.
2. Review existing Software/Service Development Lifecycle (SDLC) processes to ensure that cyber security principles are evaluated and contributed to by appropriate cyber security roles.
3. Review existing Change Management processes to ensure that cyber security principles are evaluated and contributed to by appropriate cyber security roles.
4. Develop a comprehensive incident response plan that includes layered responsibilities from the laboratory cyber office and staff with CELS and ALCF.

CSPO has committed to continuing the DSAV process in 2023. ALCF will continue to work with CSPO to verify that all Argonne security standards and practices are met.

7.2 Cyber Security Plan

Argonne's Authority To Operate (ATO) includes the ALCF as a major application and was granted on January 22, 2018. It is valid as long as Argonne maintains robust, continuous monitoring of its Cyber Security Program as detailed in the ATO letter, which is included at the end of this section.

7.3 Foreign Visitor Access and Export Controls

ALCF follows all DOE security policies and guidelines related to export controls and foreign visitor access.

Argonne is a controlled-access facility, and anyone entering the site or accessing Argonne resources remotely must be authorized. ALCF follows Argonne procedures for collecting information about foreign nationals who require site access or remote (only) computer access. All foreign nationals are required to have an active and approved ANL-593 in order to have an active ALCF account. Users can access ALCF resources only with an active ALCF account.

To apply for an ALCF account, the user fills out a secure webform in the ALCF Account and Project Management system, Userbase 3 (UB3), providing details such as legal name, a valid e-mail address, work address, phone number, and country of citizenship. They also identify the ALCF project with which they are associated. In addition, all foreign nationals (non-U.S. citizens) are required to fill out their personal, employer, demographic, and immigration/U.S. Citizenship and Immigration Services (USCIS) information in Argonne's Visitor Registration system, which is integrated with UB3. After the user submits their account application request, an e-mail is sent to the user's project PI for approval. Once the ALCF Accounts team receives the approval from the project PI, if the user is a foreign national, the user's details are electronically attached to an ANL-593 form and submitted to the Foreign Visits and Assignments (FV&A) Office for review. The FV&A Office is responsible for overseeing compliance with laboratory rules and DOE directives.

The ANL-593 form records the type of work the user will be performing, including the sensitivity of the data used and generated. The ANL-593 must be approved by Argonne Cyber Security, FV&A, the Argonne Office of Counterintelligence, and the Argonne Export Control Office. Argonne’s foreign visitor and assignments process integrates with the DOE Foreign Access Central Tracking System (FACTS), which documents and tracks access control records of international visits, assignments, and employment at DOE facilities and contractor sites. Once the ANL-593 form for the user is approved, the UB3 database is automatically updated with the user’s ANL593 start and end dates. The ALCF Accounts team then creates the user account and notifies the user. Any changes to the ANL-593 dates are automatically updated in UB3. Accounts are suspended if the user’sANL-593 expires.

ALCF allows only a limited subset of export control data on ALCF systems. ALCF works closely with Argonne’s Export Control Office to complete a detailed security plan for what export control classifications are allowed and what security measurements are required for each instance of export-controlled data. If, at any time, the ALCF wants to allow new classifications of export control data on its systems, a new security plan must be created and approved by Argonne’s Export Control Office and Argonne Cyber Security.

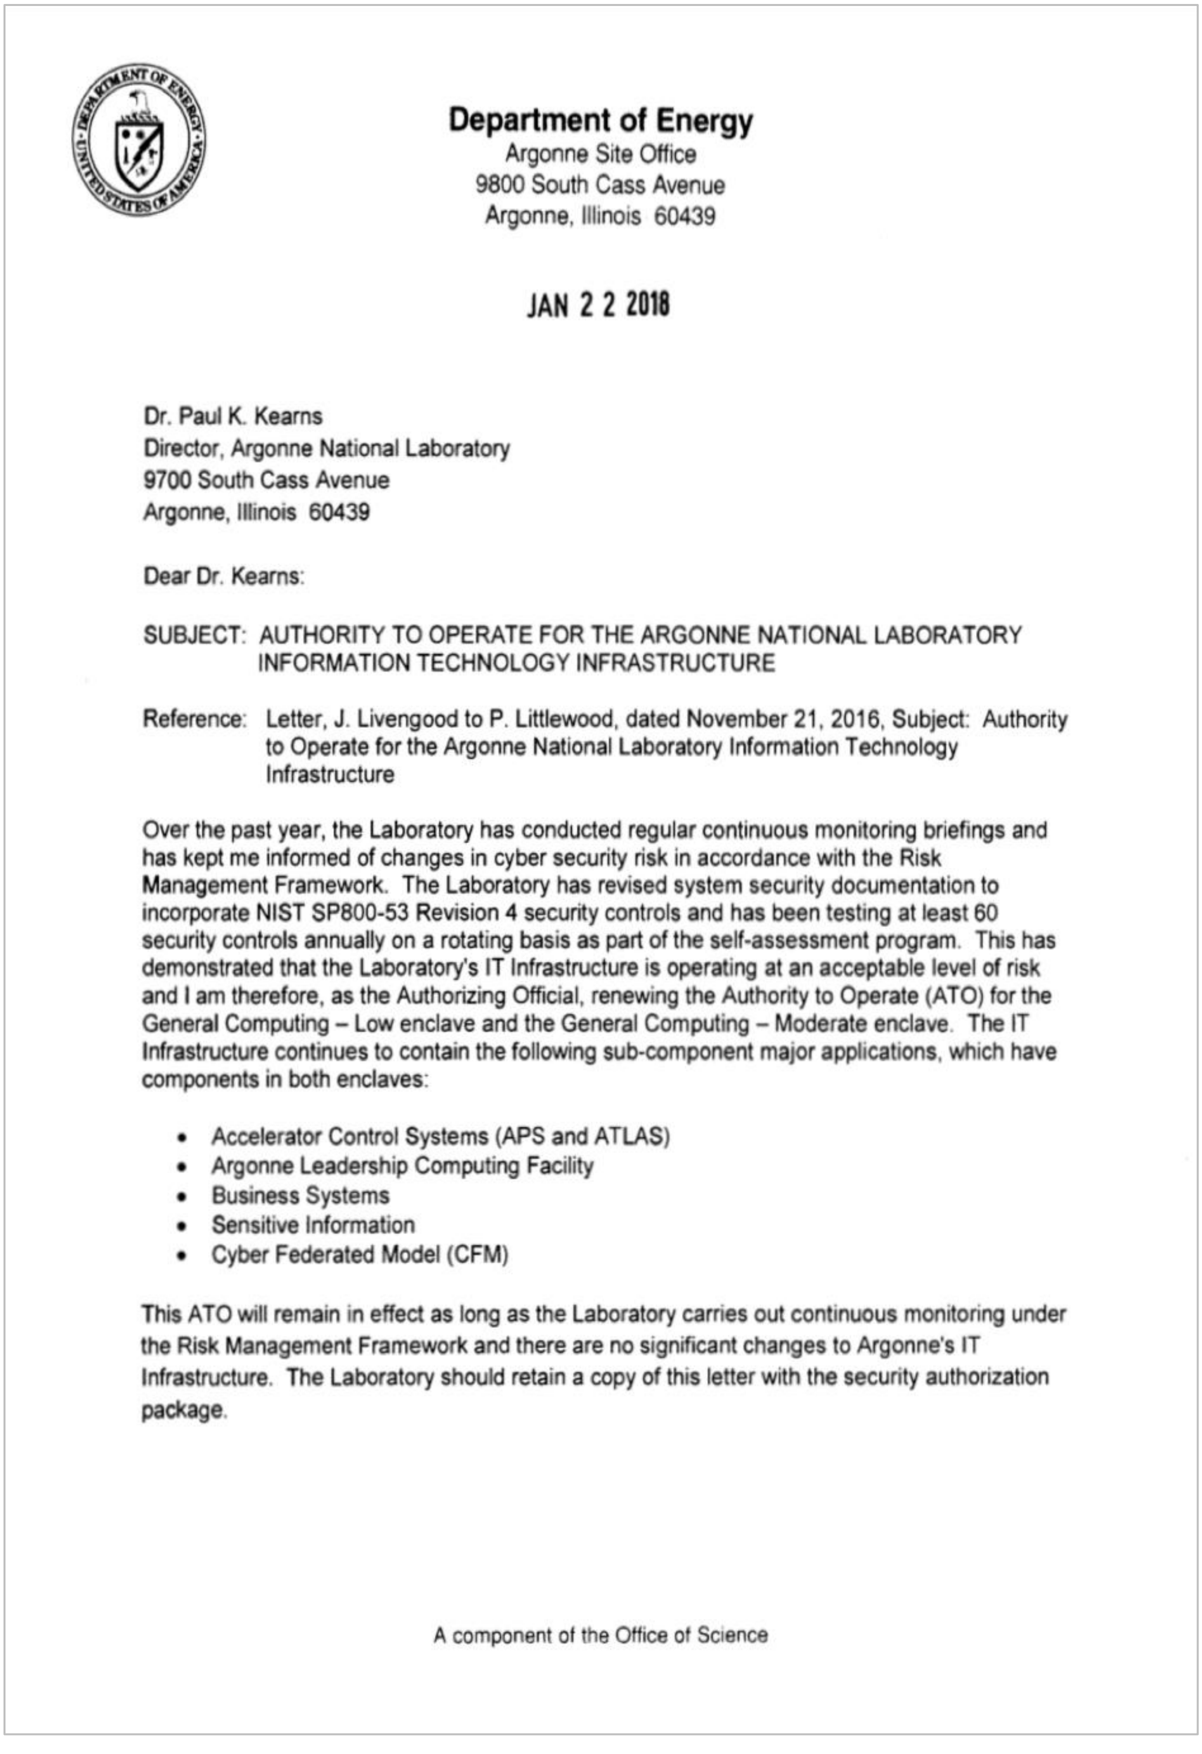

Department of Energy
Argonne Site Office
9800 South Cass Avenue
Argonne, Illinois 60439

JAN 22 2018

Dr. Paul K. Kearns
Director, Argonne National Laboratory
9700 South Cass Avenue
Argonne, Illinois 60439

Dear Dr. Kearns:

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

Reference: Letter, J. Livengood to P. Littlewood, dated November 21, 2016, Subject: Authority to Operate for the Argonne National Laboratory Information Technology Infrastructure

Over the past year, the Laboratory has conducted regular continuous monitoring briefings and has kept me informed of changes in cyber security risk in accordance with the Risk Management Framework. The Laboratory has revised system security documentation to incorporate NIST SP800-53 Revision 4 security controls and has been testing at least 60 security controls annually on a rotating basis as part of the self-assessment program. 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, renewing the Authority to Operate (ATO) for the General Computing – Low enclave and the General Computing – Moderate enclave. The IT Infrastructure continues to contain the following sub-component major applications, which have components in both enclaves:

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

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.

A component of the Office of Science

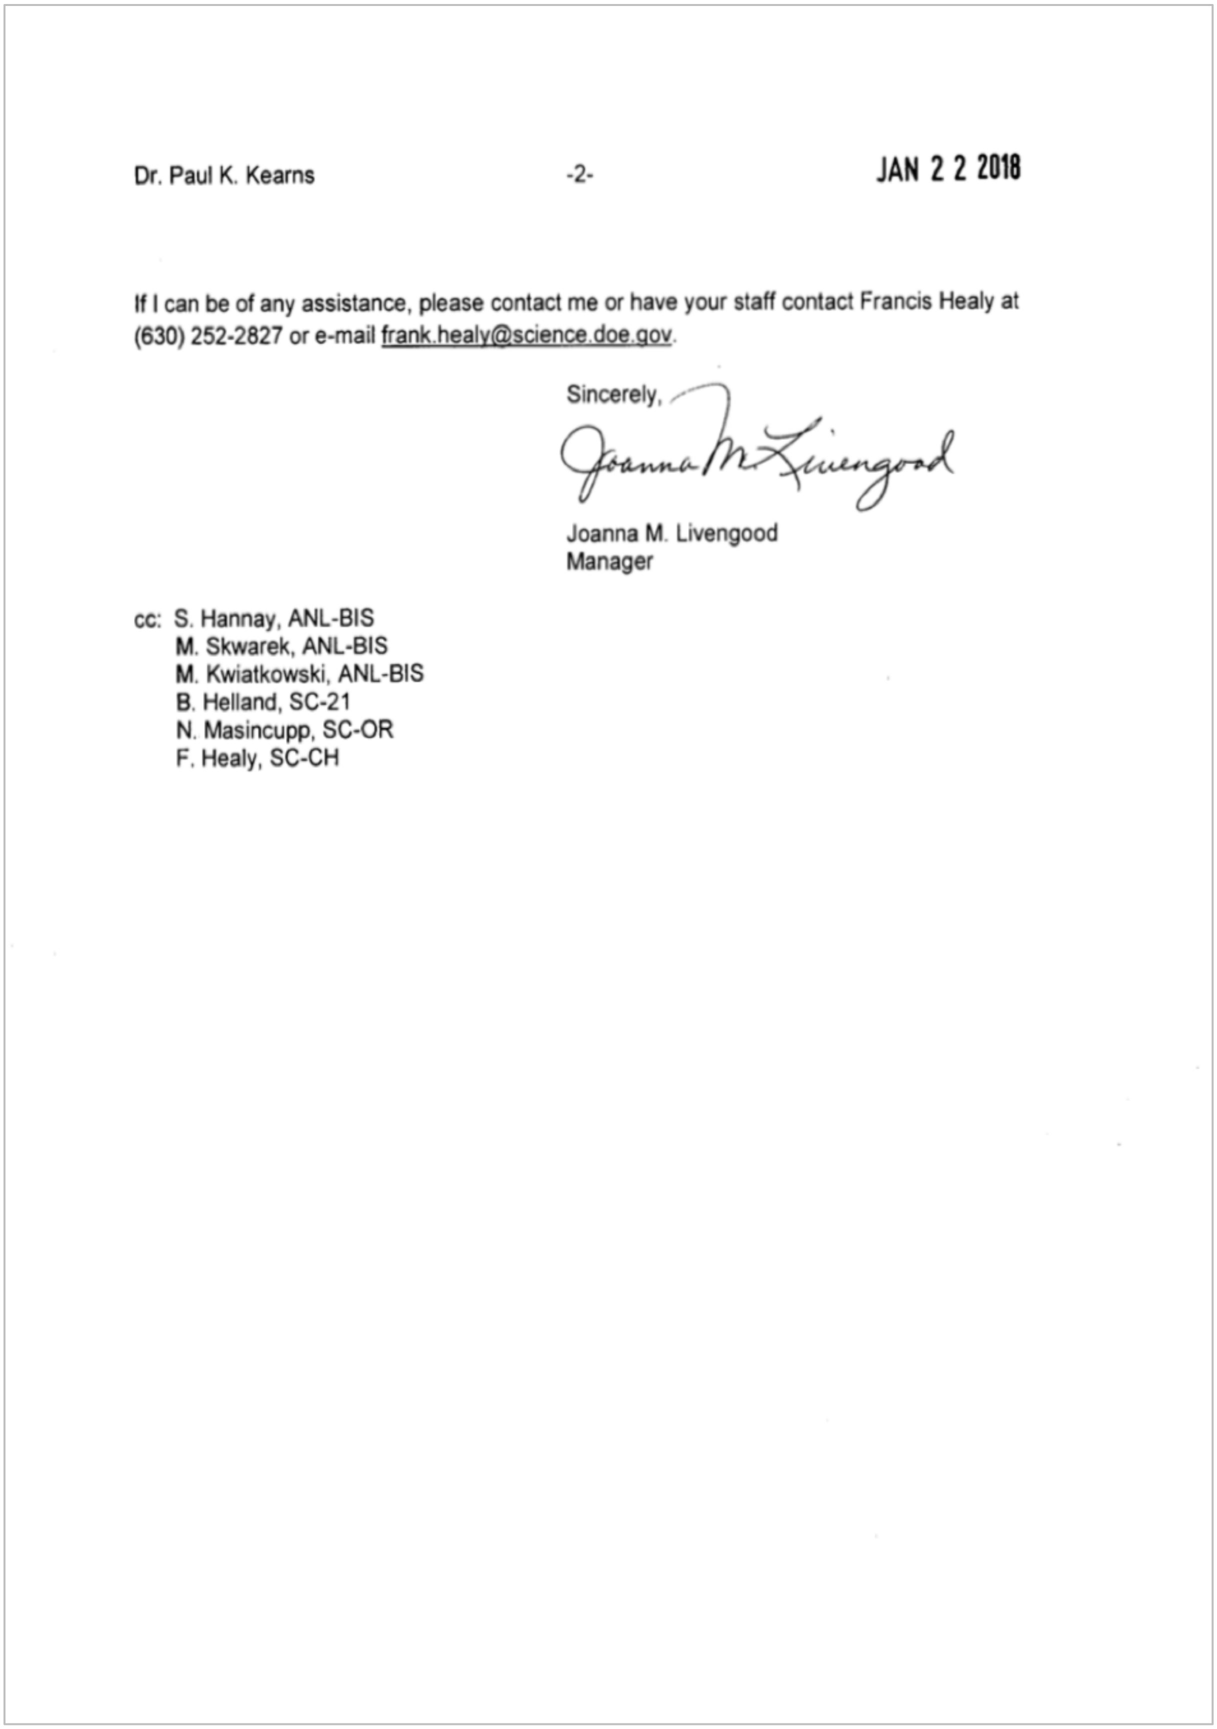

Dr. Paul K. Kearns -2- JAN 2 2 2018

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-BIS
M. Skwarek, ANL-BIS
M. Kwiatkowski, ANL-BIS
B. Helland, SC-21
N. Masincupp, SC-OR
F. Healy, SC-CH

This page intentionally left blank.

Section 8. Mission Impact, Strategic Planning, and Strategic Engagements

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

(b) Has the Facility demonstrated effective engagements with technology vendors and /or engaged in effective research that will impact next generation technology relevant to the facility's mission?

(c) Has the Facility demonstrated effective engagements with critical stakeholders (such as the SC Science Programs, DOE Programs, DOE National Laboratories, SC User Facilities, and/or other critical U.S. Government stakeholders (if applicable)) to both enable mission priorities and gain insight into future user requirements?

ALCF Response

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

In this section, ALCF reports:

- Scientific highlights and accomplishments;
- Research activities / vendor engagements for future operations; and
- Stakeholder engagement.

8.1 Science Highlights and Accomplishments

ALCF employs various methods and processes for monitoring science accomplishments. Monthly scientific highlights (mostly originating from the catalyst team) are collected and documented in a quarterly report. The determination and coordination of scientific highlights is managed by ALCF's applications team, made up of members of the catalyst team, the data science team, and the performance engineering team, and in consultation with ALCF's Director of Science. Other sources of scientific highlights include technical communications between ALCF staff members and a project PI or co-PI; significant findings reported in a high-impact publication or conference presentation; and a catalyst's own involvement in a publication.

ALCF tracks and annually reports the number of peer-reviewed publications resulting (in whole or in part) from use of the facility's resources. For ALCF, tracking takes place during a period of five years following the project's use of the facility. This may include publications in press or accepted but does not include papers submitted or in preparation. The count is a reported number, not a metric. The facility may report other publications where appropriate. Methods used for gathering publication data include asking users to verify or update ALCF's online publications database and conducting Google Scholar and Crossref searches. ALCF also collects approximately one-third of its users' ORCID iDs in any given year and has been investigating ways to use this method to identify more user publications.

Table 8.1 shows the breakdown of refereed publications based, in whole or in part, on the use of ALCF resources, and highlights those appearing in major journals and proceedings. These include three publications in *Nature Communications*, two in *npj Computational Materials*, one in *Nature Physics*, *two* in *Scientific Reports*, four in *Scientific Data*, one in *npj Climate and Atmospheric Science*, one in *Nature Catalysis*, and one in *Communications Engineering* (combined in the *Nature* journals category in the table below); one in *Science Advances* (listed under the *Science* journals category in the table below); one in the *Proceedings of the National Academy of Sciences (PNAS)*; three in *Physical Review Letters*; and six in the proceedings of the *2022 International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)*. Table 8.2 shows updated publication counts from prior years, based on new information received after the prior year's OAR deadline.

Projects using ALCF resources earned multiple awards in 2022, including three awards from HPCwire, a leading website covering the HPC community. The HPCwire Readers' Choice Award for Best Use of High-Performance Data Analytics * Artificial Intelligence category was given to an Argonne-led multi-institutional team that showcased how to create and share FAIR (Findable, Accessible, Interoperable, and Reusable) data and AI models within a unified computational framework and how it may be harnessed to enable autonomous AI-driven discovery. Argonne together with Raytheon Technologies Research Center won a Readers' Choice Award for Best Use of HPC in Industry *(Automotive, Aerospace, Manufacturing, Chemical)* for developing machine learning models for designing and optimizing high-efficiency gas turbines in aircraft. In the same industry category, Argonne also won an Editors' Choice Award for its work with Aramco Americas and Convergent Science focused on high-fidelity CFD simulations of hydrogen engines using resources at ALCF and Argonne's Laboratory Computing Resource Center. An Argonne-led team won yet another prestigious award at SC22—the *ACM Gordon Bell Special Prize for HPC-based COVID-19 Research* for their new method of quickly identifying how a virus evolves, which involved significant support and engagement from ALCF staff members and many others across Argonne. The winning team's work in training large language models (LLMs) to discover variants of SARS-CoV-2 has implications for biology beyond COVID-19.

Table 8.1 Summary of Users' Peer-Reviewed Publications in CY 2022

Nature Journals	Science	PNAS	Physical Review Letters	SC	Total 2022 Publications
15	1	1	3	6	213

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

OAR Year	CY 2018	CY 2019	CY 2020	CY 2021	CY 2022
Total Publications	276	288	257	249	213

Science Highlights

Scientific highlights are short narratives that illustrate the user facility's contribution to advancing DOE strategic goals. Highlights may describe a research accomplishment or significant finding from either a current project or from a project originating in a previous year, as data analysis may occur several months after the computational campaign has been completed.

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 (in parentheses) shows 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," and "high capability" from left to right.

8.1.1 Computer Simulations Aided Design of Environmentally Benign Electrolyte for Lithium Battery

The Science

For scientists working to create the next generation of batteries, water is a formidable enemy. If a lithium-ion battery contains any water, it won't work in the conditions it needs to retain stability and safety. The challenge is rational design of an electrolyte that is immune to water with reduced hazardous risks. A joint experimental and computational approach used MD (Molecular Dynamics) simulations to model the role of water under realistic experimental conditions and identified the mechanism by which a novel electrolyte can literally hold water against conventional "dry" electrolytes. Enhanced sampling of molecular simulations provided insight into water clustering, ion distribution, and structure of the cathode-electrolyte interface that gave rise to the favorable battery performance.

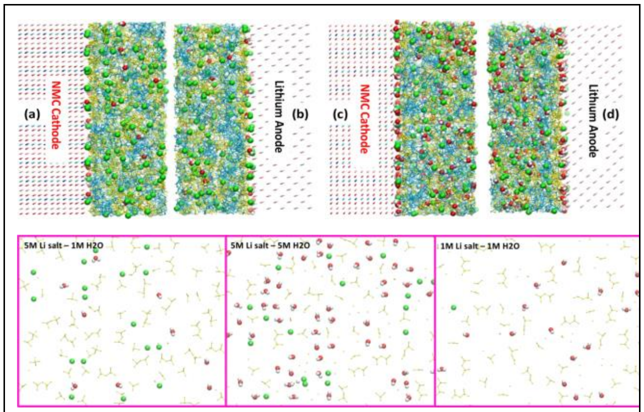

Figure 8.1 Simulated cluster of water (red) and lithium ions (green) at electrode surfaces under various water concentrations. Battery performance is retained when water molecules remain isolated.

Image courtesy of Wei Jiang (Argonne) and reproduced from Figures 6 and 7 of *ACS Appl. Mater. Interfaces*, **13**, 58229 (2021).

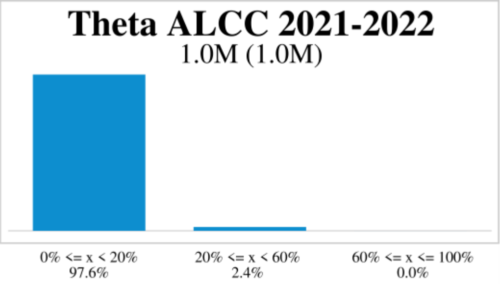

The Impact

In state-of-the-art battery manufacturing, the required critical moisture controls involve preparing battery components in a dry environment — a tremendously energy-intensive effort. An electrolyte's immunity to moisture in the environment would thus eliminate this stringent control for electrolyte formulation, storage, and transportation. It would also facilitate mass production and reduce battery cost, given the reduced need for energy to remove residual water from battery components. To this end, researchers have demonstrated a novel electrolyte which can accommodate a thousand times as much water as conventional electrolytes and shows exceptional stability on general battery electrodes.

Summary

The ensemble-based Hamiltonian Annealing (HA) method is a powerful sampling enhancement algorithm that accelerates infrequent configurational transitions of complex liquids by annealing the interaction energy of a simulated system. The HA simulations were critical to searching all important water configurations in a viscous electrolyte medium that would otherwise have required inaccessible timescales.

The team processed the MD trajectory data using in-house configuration analysis software to quantify the water clustering within the electrolyte and at the electrode surfaces (Figure 8.1). Isolated water molecules are regarded as friendly additives, whereas any continuous water clusters can cause instability and even hazardous reactions. At a specific range of water concentration ratios in the salty medium, water molecules were completely sequestered by electrolytes and lithium ions, and no water clustering was observed. The statistical analysis revealed the chemical mechanism that only specific molecular structures and compositions of electrolyte components can sequester water molecules. The simulated “water sequestering” phenomena guided battery fabrication experiments where electrolytes were exposed to environmental moisture, and these experiments demonstrated the stability of a lithium battery with the humidified electrolyte. The joint computational and experimental approach can offer a pathway toward moisture-immune electrolyte design for future lithium battery fabrication.

ALCF Contribution: ALCF staff built the NAMD code, helped maintain optimal usage of the resources by coordinating the running of production jobs, and managed the large amount of simulation data. The ASCR allocation PI, Wei Jiang, developed and efficiently leveraged the ensemble sampling enhancement method.

Contact

Zhengcheng Zhang
Argonne National Laboratory
zzhang@anl.gov

ASCR Allocation PI: Wei Jiang, Argonne National Laboratory

Publication

Q. Liu, W. Jiang, Z. Yang, and Z. Zhang, “An Environmentally Benign Electrolyte for High Energy Lithium Metal Batteries,” *ACS Appl. Mater. Interfaces*, **13**, 58229 (2021). DOI: 10.1021/acsami.1c19124

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR: June 6, 2022

8.1.2 Metastable Phase Diagrams for Materials

The Science

Phase diagrams are an invaluable tool for materials synthesis, providing researchers with information on the phases of a material at any given thermodynamic condition (e.g., pressure, temperature, chemical composition). Materials may not reach their equilibrium state during synthesis, operation, or processing and may remain trapped in a local (metastable) free energy minimum. To extend their utility to a promising but mysterious class of materials, Argonne researchers are developing an automated workflow to construct phase diagrams for metastable materials using machine learning. The workflow successfully predicted the equilibrium phase for carbon and several observed and predicted metastable structures (Figure 8.2). The workflow also identified a new metastable phase of carbon that helps to resolve prior experimental observations.

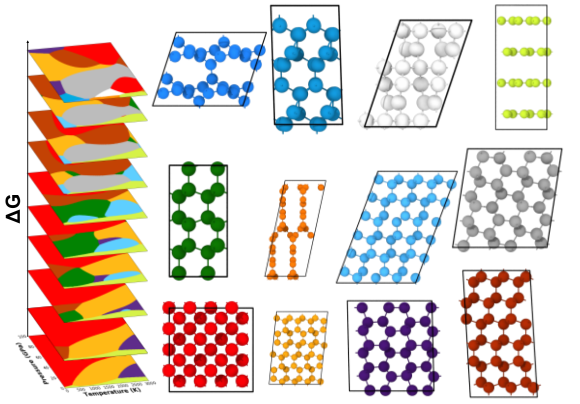

Figure 8.2 The final product of the machine learning algorithm: metastable phase diagrams for carbon. The colored regions indicated conditions at which carbon exists in certain metastable states (with similarly colored structures) that may yield useful material properties.

Image courtesy of Argonne National Laboratory.

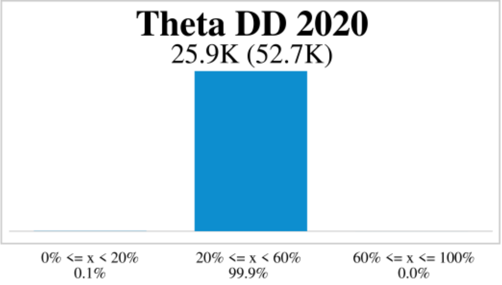

The Impact

Mapping metastable phases and their thermodynamic behavior is a highly desirable but nontrivial and data-intensive task due to the vast configurational landscape. The team's automated framework for constructing metastable phase diagrams lays the groundwork for computer-aided discovery and design of synthesizable metastable materials, which could help advance a range of applications including in semiconductors, catalysts, and solar cells. Metastable phase diagrams not only help accelerate phase identification by narrowing the list of potential candidate structures, but more importantly aid in discovering novel polymorphs.

Summary

Thermodynamic phase diagrams provide information about stable states as a function of various intensive properties (e.g., temperature, pressure, and chemical composition). The first challenge in predicting phase diagrams for metastable materials is to efficiently identify the global and local minima of the free energy surface. The next is mapping the free energy for each identified phase over a range of thermodynamic variables to determine regions of relative stability. The last is to classify and identify phase boundaries and domains of metastability for the structures.

Initial pools of candidate structures were constructed using (1) genetic algorithm and geometry optimizations computed using DFT (Density Functional Theory) with the VASP code, and (2) evolutionary structure search using the LCBOP forcefield and LAMMPS. All structures identified within a specified enthalpy cutoff were selected for further analysis. Geometrically similar structures were placed into groups based on layered structures, radial distribution functions, and angular distribution functions. The structures in each of the 505 groups with the lowest enthalpy were selected for free energy calculations. A surrogate machine learning model was then trained to predict free energy as a function of temperature and pressure using 273 of the candidate phases. Additional structures from the SACADA database were evaluated for inclusion with the final metastable phase diagram consisting of 18 phases. Additional analyses were performed once regions of metastability were identified. Transformation barriers between metastable phases were estimated using solid-state nudged elastic band calculations and the VASP code. Domains of synthesizability were estimated from LAMMPS simulations with the LCBOP model and testing for structure deformation at different temperatures and pressures.

ALCF Contribution: ALCF staff assisted with LAMMPS on Theta and with scheduling jobs to meet a publication deadline.

Contact

Subramanian Sankaranarayanan
Argonne National Laboratory
ssankaranarayanan@anl.gov

Publication

S. Srinivasan et al., “Machine learning the metastable phase diagram of covalently bonded carbon,” *Nat. Commun.*, **13**, 3251 (2022). DOI: 10.1038/s41467-022-30820-8

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR: September 14, 2022

8.1.3 Next-Generation Nonwovens Manufacturing: A Model-Driven Simulation and Machine Learning Approach

The Science

The team developed a computational fluid dynamics (CFD) model using large eddy simulations and volume-of-fluids (VOF) to understand air flow behavior and liquid/gas interface dynamics, which enabled exploring varying geometry and operating conditions (selected using active machine learning [ML] approaches) with the CFD model. As a result, the team identified process-relevant success metrics using the collected data, employing a Bayesian optimization strategy (with a Gaussian process surrogate model) to optimize the process and iteratively recommend simulation parameters to test and thereby reach a desired configuration with reduced energy consumption.

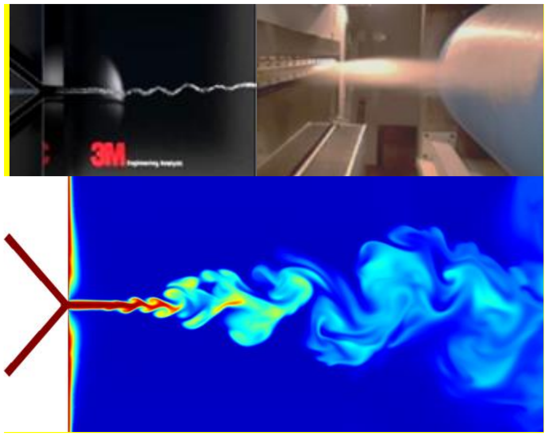

Figure 8.3 (top) Melt-blowing experiments at 3M and (bottom) Argonne CFD simulation results modeling the same process using conditions selected from machine learning.

Image courtesy of the Argonne and 3M team members on the project.

The Impact

Nonwoven materials prepared via melt blowing are widely used to make filters (e.g., N95 masks), fabrics, and insulation materials (Figure 8.3). This project aims to reduce the process energy consumption without compromising quality through a coordinated campaign of experiments, simulations, and ML. The goals are to optimize energy efficiency for nonwoven applications — a 300,000-ton materials market — while maintaining process quality. Reducing overall energy consumption by 20% would mean a global impact for 3M, and other manufacturers would likely follow suit.

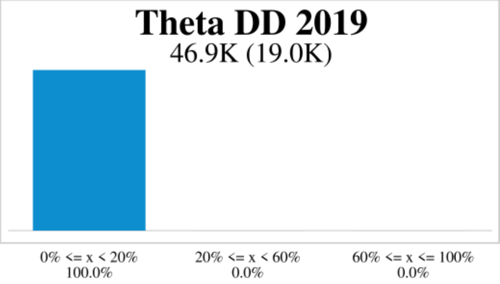

Summary

After completing several CFD simulations using CONVERGE on Theta to model this process, several classical ML methods were integrated to predict ideal fluid flow patterns at untested process conditions and apparatus geometries. The objective was then to predict novel geometries that will minimize the air channel angle (ACA), a critical performance parameter.

To enable the ACA predictions, a random forest regressor and a multiple linear regressor with a leave-one-out validation approach were trained to predict the air cone angle from the geometry and inlet conditions. It was demonstrated via the leave-one-out method that ACA can be predicted to within a root mean square error (RMSE) of roughly 3 degrees. To minimize the ACA, inlet conditions that satisfy this criterion were then predicted. An optimal process condition to check on was found to have a blade angle of 56.9 degrees, inlet temperature of

505.6 Kelvin, and inlet pressure of 113.4 KPa, resulting in a channel angle of 32.014 +/- 0.341 degrees and possibly providing significant energy savings.

ALCF Contribution: ALCF staff helped the team early on with understanding CFD benchmarks with OpenFOAM on Theta and by providing guidance on planning job submissions and DD allocations.

Contact

Ian Foster
Argonne National Laboratory
foster@anl.gov

Publications

B. Bass, J. R. New, and W. Copeland, "Potential energy, demand, emissions, and cost savings distributions for buildings in a utility's service area," *Energies* **14**, 132 (2021). DOI: https://doi.org/10.3390/en14010132

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR: February 16, 2022

8.1.4 Extreme-Scale Simulations for Advanced Seismic Ground Motion and Hazard Modeling

The Science

The creation and dynamics of surface ruptures during earthquakes depend on many factors, such as proximity to the fault system, fault geometry, and surrounding material properties (e.g., soil and rock). Simplified models can be helpful for performing rapid seismic hazard analyses to inform urban planning and building regulations. However, empirical fault-displacement models are sparse and poorly constrained due, in part, to a lack of high-resolution observations. This project used observations from the 1992 Landers earthquake, the third-largest California event of the 20th century with a magnitude of 7.3, to calibrate and validate a dynamic rupture model that reproduced several first-order fault-displacement metrics (Figure 8.4), such as the location of peak displacement.

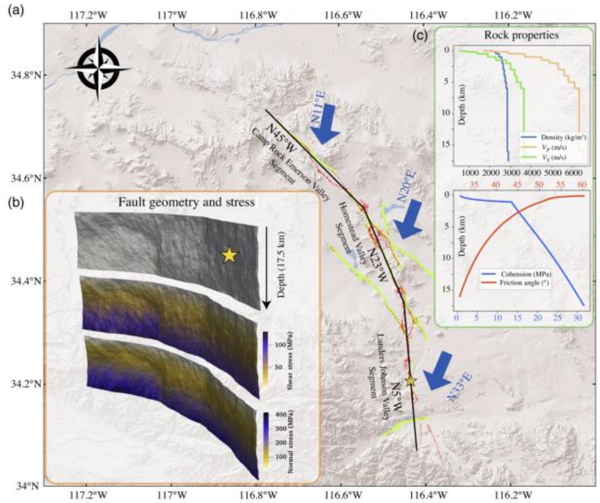

Figure 8.4 Simulation model setup showing (a) observational and model fault details with arrows indicating orientations of principal stresses, (b) fault-plane model and epicenter (star), and (c) rock properties.

Image courtesy of Yongfei Wang and Christine Goulet, USC.

The Impact

Fault displacements in large earthquakes have caused significant damage to structures and lifelines and represent a real threat to distributed infrastructure systems spanning faults in more than one location. Physics-based dynamic rupture models and simulations that can capture the general displacement characteristics observed, such as in this work, serve as useful tools both for extrapolating to new scenarios and informing the development of probabilistic fault-displacement hazard analyses. Improved hazard assessments can better inform and prepare residents for the hazard of earthquakes, enabling strategies to mitigate societal and economic impacts and to save lives in the event of a major earthquake.

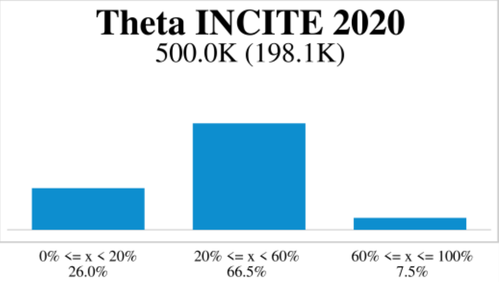

Summary

Observational data for constructing robust empirical fault-displacement models are scarce due to (1) low occurrence of earthquake ruptures that reach the surface (such large-magnitude events tend to be rare); (2) a long recurrence period between earthquakes, preventing comparison of potentially similar events; and (3) the technical difficulty of measuring fault displacements over very large areas. The rupture dynamics model, while it does not reproduce every single displacement observed in a specific event such as the Landers case, does capture the general displacement characteristics observed from an ensemble of simulations.

Three key ingredients in the model are (1) fault geometry and initial stress condition, (2) fault failure criteria indicating when slip is allowed, and (3) the surrounding material (rock) properties. The fairly complex fault systems involved during Landers were modeled as three linear segments loosely fitting the key faults with added roughness perturbations. The preferred model was parameterized to be the most representative of the as-observed Landers event and reproduced multiple first-order fault-displacement metrics. Derived from this model, a suite of ensemble models was defined spanning four sets of physical parameters and 12 realizations of the fault roughness. The Support Operator Rupture Dynamics (SORD) code was used to simulate the fault displacements by numerically solving the 3D elastoplastic spontaneous rupture propagation problem. The model contains 1.9 billion hexahedral elements and simulates a 60 s rupture. While the simplified model was not able to capture small-scale fault displacements or complex features observed in the actual event (e.g., the multi-segment rupture pattern), it did reproduce several intermediate- to large-scale features, such as (1) total displacement, (2) mean off-fault deformation ratio, (3) mean fault-zone width, and (4) location of the peak displacement. Both Frontera at TACC (Texas Advanced Computing Center) and Theta at ALCF were used in this work.

ALCF Contribution: ALCF staff helped the team with debugging initial issues on Theta: scheduling 4096-node jobs for Waveqlab3D and compiling and running SORD with Cray MPI.

Contact

Yongfei Wang and Christine Goulet
University of Southern California
yongfeiw@usc.edu; cgoulet@usc.edu

ASCR Allocation PI: Christine Goulet, University of Southern California

Publications

Y. Wang and C. Goulet, “Validation of Fault Displacements from Dynamic Rupture Simulations against the Observations from the 1992 Landers Earthquake,” *Bull. Seismol. Soc. Am.*, **111,** 2574-2494 (2021). DOI: 10.1785/0120210082

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR: May 9, 2022

8.1.5 Binding Energies from Effective Field Theory with Quantified Uncertainties

The Science

Developing a reliable, quantitative first-principles description of nuclear structure and reactions with quantified uncertainties remains one of the main challenges in computational nuclear physics. Presently, the most promising approach to reach this ambitious goal combines chiral effective field theory (EFT) to describe nuclear interactions in harmony with the symmetries (and their breaking pattern) of QCD with *ab initio* few-body methods to tackle the quantum mechanical A-body problem. To address this challenge, the Low Energy Nuclear Physics International Collaboration (LENPIC) aims to develop accurate and precise two- and three-nucleon interactions by pushing the EFT expansion to high chiral orders and using these interactions to solve the structure and reactions of light nuclei.

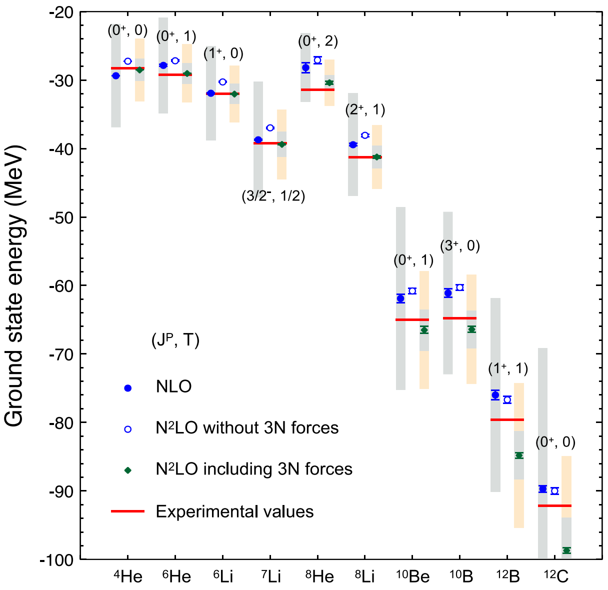

Figure 8.5 Calculated ground state energies in MeV using chiral interactions in comparison with experimental values.

Image courtesy of Evgeny Epelbaum.

The Impact

Using up to full-machine runs on Mira and Theta, the team performed the first tests of novel chiral nucleon-nucleon potentials with consistent three-nucleon interactions. This result demonstrates the importance of three-nucleon interactions and allows for a quantitative understanding of the theoretical uncertainties due to the chiral EFT expansion. The team also extended and tested a Bayesian statistical model that learns from the order-by-order EFT convergence pattern to account for correlated excitations. This step allowed the team to demonstrate agreement with experimental ground state energies as well as excitation energies to within the estimated theoretical uncertainties (Figure 8.5).

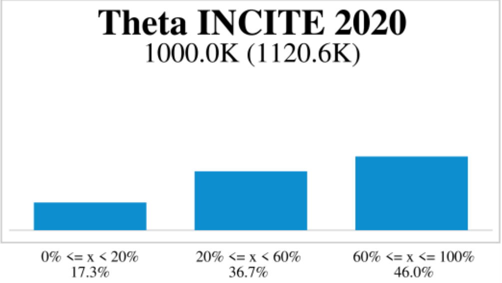

Summary

The team used the No-Core Configuration Interaction (NCCI) approach as implemented in their Many Fermion Dynamics nucleon (MFDn) code for the calculations. Most runs were done on Mira, with some additional runs on Theta and Cori. The configuration of Theta, in combination with the policy that users can request the KNL nodes either in quad-cache mode or in quad-flat mode, was very useful, as the team's code is memory bound. For the largest runs, the team used almost the entire Theta machine in quad-flat mode, which gave the best performance and maximal memory.

These large-scale runs allowed the properties of the ground and excited states of light nuclei to be calculated with robust theoretical error estimates. These results were then compared to known experimental results to test consistent LENPIC chiral EFT interactions with 2- and 3-nucleon interactions. The calculated results were consistent with the experimental results, confirming the validity of the approach.

ALCF Contribution: ALCF staff helped with performance optimization for Mira, and with debugging issues and job scheduling on both Mira and Theta. ALCF also provided excellent training opportunities for porting and performance tuning on both Mira and Theta. The training for Mira was essential for this work, and the training on Theta was useful for understanding the differences between Theta and Cori.

Contact

Pieter Maris
Iowa State University
pmaris@iastate.edu

ASCR Allocation PI: Gaute Hagen, Oak Ridge National Laboratory

Publications

P. Maris et al., *Phys. Rev. C*, **103**, 054001 (2021).

Highlight Categories

Performer/Facility: ASCR-ALCF

Date Submitted to ASCR: March 2022

8.2 Research Activities / Vendor Engagements for Future Operations

8.2.1 Research Activity - Joint Laboratory for System Evaluation (JLSE)

Argonne's JLSE enables researchers to assess and improve next generation computing platforms of interest to the DOE. Established by the computing divisions of Argonne's CELS Directorate (Data Science and Learning, Mathematics and Computer Science, Computational Science, and ALCF), and run by ALCF, 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 systems-level experimental computer and computational science, including operating systems, messaging, compilers, benchmarking, power measurements, Input/Output (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's and DOE's needs in a variety of areas 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 exascale 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.

ALCF closely collaborates with Intel on Aurora. This includes accelerating their software roadmap for traditional HPC and for data and AI pillars to support the science workloads from ESP and ECP projects. JLSE testbeds and software used to prepare for Aurora include:

- **Florentia:** Six nodes with Intel Server Board (codenamed Denali Pass), and four Intel Data Center GPU Max (codenamed Ponte Vecchio) and two Xeon CPU Max Processors (codenamed Sapphire Rapids) in each node.
- **Arcticus:** Seventeen nodes with Intel Server Board (codenamed Coyote Pass) with two Intel development GPU cards (codenamed XeHP_SDV) and two Intel Xeon Gold 6336Y CPUs in each node. A similar cluster (called DevEP) with 32 nodes at Intel was also made available to JLSE users.
- **Iris:** SuperMicro X11SSH-GF-1585 Server Motherboard with Intel Xeon E3-1585 v5 CPU and Iris Pro Graphics P580 GPU (Intel integrated Gen9 GPUs).
- **Intel Pre-production Development Platform** includes 2x Next Gen Intel Xeon Scalable processor.
- **Presque:** Intel DAOS nodes (DCPMM and NVMe storage) with Intel DAOS file system.

Other JLSE active testbeds include:

- Intel Xeon Phi (codenamed Knights Landing) Cluster

- NVIDIA GPUs:
    - DGX–1 (V100 GPUs)
    - Gigabyte NVIDIA A100 and A40 cluster
    - Supermicro NVIDIA V100 and P100 cluster
- AMD GPUs:
    - AMD GPU MI250, MI100 and MI50 cluster
- Intel Xeon Clusters: Skylake, Cascade Lake, and Cooper Lake
- ARM Clusters:
    - HPE Apollo 70 – Comanche Prototype ARM64 Cluster
    - HPE Apollo 80 – 8 node Fujitsu A64X CPU
    - NVIDIA ARM Dev Kit – Ampere Altra Q80-30 ARM CPU, NVIDIA A100 GPU
- IBM Power System AC922 (Power9 CPU, V100 GPU)
- Atos Quantum Learning Machine

In 2022, the JLSE supported more than 500 users spanning more than 150 projects. These projects ranged from application portability to software development—including operating systems, compilers, deep learning frameworks, and performance tools. Teams from within the ECP's Application Development and Software Technology groups have been using the JLSE Aurora testbeds and the Aurora SDK to develop applications and software for Aurora. The following summaries represent a sampling of current JLSE projects:

- **Exascale Computing Project:** Projects from ECP were given access to the Intel GPU nodes in Florentia and Arcticus and oneAPI software to port their applications and software to the platform.
- **ALCF Early Science for Aurora:** ESP application teams used JLSE resources to prepare and optimize applications for the next-generation supercomputers in advance of Aurora becoming available. For example, researchers from the Aurora ESP projects access the Xeon Skylake Iris nodes with Intel's integrated GPUs and the early versions of oneAPI software to develop and test their applications for Aurora.
- **Big Data:** Researchers used JLSE 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 (using Java and Scala) that can run on a supercomputer, cluster, or a cloud with a back end for executing data-intensive communication patterns.
- **Deep Learning:** Multiple projects used JLSE systems to investigate the potential of deep learning. One research team focused on understanding how deep learning can be used to improve lossy compression of scientific data from simulations and instruments. Another team explored the performance of different machine learning frameworks that have implemented deep learning and neural networks on KNL systems.
- **Compilers:** The JLSE testbed was used to verify the latest version and new features of LLVM on a variety of architectures that are not the researcher's normal development environment and then to measure performance changes. Specifically, the continued development of the LLVM OpenMP runtime used the JLSE machines with NVIDIA, AMD, and Intel accelerators to ensure cross-platform compatibility and identify problems early. The IBM PowerPC system can be used as a stand-in for Summit without providing access to it (e.g., for students). Also, compilation of LLVM itself can complete within

minutes (instead of an hour on consumer-grade hardware), allowing for a faster development cycle. A nightly build of the latest LLVM version was made available to all JLSE users so they could test their applications and report compiler bugs before reaching an official release version.

- **MPI:** Several improvements to the MPICH implementation of MPI were tested on JLSE systems, including the GPU-aware communication utilizing GPU IPC and GPU DirectRDMA, enhanced threading support through implicit and explicit communication context mapping, and GPU-stream-triggered MPI operation, which allows MPI operations to be enqueued and initiated by GPU.
- **Quantum Computing:** A research team used the JLSE's Atos Quantum Learning Machine and other quantum simulators to develop and apply quantum algorithms like variational quantum eigensolver (VQE). Such algorithms are used to improve the quality of quantum chemistry calculations on quantum computers using a hybrid approach. In particular, the simulators are used to develop low-depth quantum circuits for modern quantum computers with the end goal of accurately calculating enthalpies of formation, atomization energies, ionization energies, and electron affinities.
- **SYCL Programming Model:** ALCF continued to work with Codeplay, NERSC, and OLCF in porting SYCL to NVIDIA A100 GPU and AMD GPU resources in JLSE. This would allow SYCL-based applications written for Aurora to run on Polaris, Perlmutter, Frontier, and other systems based on NVIDIA and AMD GPUs.

8.2.2 Research Activity - ALCF AI Testbed

With an eye toward the future of scientific computing, ALCF has deployed an advanced AI platforms testbed for the research community. This testbed enables the facility and its user community to help define the role of AI accelerators in next-generation scientific machine learning. It also helps shape the roadmap and development of AI accelerators for science. The testbed's innovative AI platforms complement Argonne's GPU-accelerated supercomputers, Polaris and Aurora, to provide a state-of-the-art computing environment that supports pioneering research at the intersection of AI and HPC.

The ALCF AI testbed consists of systems from Cerebras, Graphcore, Groq, Intel Habana, and SambaNova. ALCF actively works with several AI accelerator systems and plans to include new systems as part of the testbed. The Cerebras and SambaNova systems were made available to the open-science user community and can be requested as part the DD allocation program. Active users of the systems span university, industry, and national labs and include applications in domains such as material science, cosmology, bioscience, imaging science, high energy physics, and climate sciences. ALCF expects to make more systems available in 2023. Three highlights in 2022 are the ACM Gordon Bell Special Prize for COVID-19 (SC 2022); use of the Groq system to help improve the operations of future fusion energy devices; and a comprehensive evaluation of various AI accelerators for scientific machine learning (PMBS22). ALCF will upgrade the CS-2 system to include an appliance mode to facilitate larger-scale models. The SambaNova system will be upgraded to the new 2nd-generation accelerator, scaled out to eight nodes. ALCF also procured a rack-scale Graphcore Bow-64 system. ALCF has conducted several user workshops on the AI testbed and held a tutorial at SC22 in collaboration with AI testbed partners to help the community leverage these systems for science.

The ALCF AI Testbed includes the following systems:

- The **Cerebras** CS-2 is a wafer-scale deep learning accelerator comprising 850,000 processing cores, each providing 48KB of dedicated SRAM memory for an on-chip total of 40 GB interconnected to optimize bandwidth and latency. Its software platform integrates popular machine learning frameworks such as TensorFlow and PyTorch.
- The **SambaNova** DataScale system is architected around the next-generation Reconfigurable Dataflow Unit (RDU) processor for optimal dataflow processing and acceleration. The SambaNova is a half-rack system consisting of two nodes, each of which features eight RDUs interconnected to enable model and data parallelism. SambaFlow, its software stack, extracts, optimizes, and maps dataflow graphs to the RDUs from standard machine learning frameworks, including TensorFlow and PyTorch.
- The **Graphcore** 22 petaflops Bow Pod64 system is the latest-generation accelerator from Graphcore. This is a one-rack system consisting of 64 Bow-class IPUs with a custom interconnect. The Graphcore software stack includes support for TensorFlow and PyTorch. It includes the Poplar SDK used by machine learning frameworks.
- The **Habana** Gaudi processor features eight fully programmable VLIW SIMD tensor processor cores, integrating ten 100 GbE ports of RDMA over Converged Ethernet (RoCE) into each processor chip to efficiently scale training. The Gaudi system consists of two HLS-1H nodes, each with four Gaudi HL-205 cards. The software stack comprises the SynapseAI stack and provides support for TensorFlow and PyTorch.
- A **Groq** Tensor Streaming Processor (TSP) provides a scalable, programmable processing core and memory building block able to achieve 250 TFlops in FP16 and 1 PetaOp/s in INT8 performance. The Groq accelerators are PCIe gen4-based, and multiple accelerators on a single node can be interconnected via a proprietary chip-to-chip interconnect to enable larger models and data parallelism.

Key activities of the testbed include:

- Maintaining a range of hardware and software environments for AI accelerators.
- Providing a platform to benchmark applications, programming models, and ML frameworks.
- Supporting science application teams in the porting and evaluation of their applications.
- Coordinating with vendors during their product development.

The AI Testbed effort supports remote access to the systems, collects feedback and use cases from users, develops online tutorials in conjunction with each of the vendors, and conducts in-person training and hackathon events.

Common Software Environment: ALCF worked with AI testbed vendors to use PBSPro to manage and schedule resources. This will enable us to better integrate the AI testbed with the rest of the ALCF resource complex.

8.2.3 Vendor Engagement – Codeplay Software Ltd. (Codeplay)

ALCF collaborated closely with Codeplay Software, OLCF, and NERSC in continuing support of the SYCL 2020 programming model on both AMD and NVIDIA GPUs. Codeplay creates software based on open programming standards so that application developers can program complex processors using familiar standards and tools. In 2022, Codeplay Software released two binary packages, one for NVIDIA GPUs and one for AMD GPUs, which integrate with the binary packages from Intel's oneAPI suite. This allows for any systems running NVIDIA or AMD GPUs to use Intel's oneAPI, and users have a portable accelerator programming environment that's compatible with all three types of accelerator hardware. This work was enabled by all four collaborators.

8.2.4 Vendor Engagement – Altair and OpenPBS

The ALCF team continued its collaboration with Altair Engineering and the OpenPBS community. In 2022, ALCF completed the work on eliminating the Python Interpreter memory leaks. There is one related issue to fix, which is the Interpreter restart; however, the priority has shifted to porting PBS to the AI testbed system and supporting Sunspot and Aurora. ALCF also identified and worked with Altair to resolve five bugs in 2022. ALCF was periodically briefed on and provided input into Altair roadmap items, which included the upgrade to Python 3.9, GraphQL, and multi-server during CY 2022. Finally, due to internal process changes, Altair was not pushing appropriate fixes out to the OpenPBS community. It was resolved once ALCF raised that as an issue.

8.2.5 Vendor Engagement – DOE Advanced Computing Ecosystem RFI

DOE's Office of Science and National Nuclear Security Administration asked six laboratories—Argonne, Berkeley, Oak Ridge, Livermore, Los Alamos, and Sandia—to work together to engage the computing technology and systems vendor community to gathering information on the future technologies and products that could be relevant to advanced computing ecosystem requirements, specifically those for computational and data science approaches and solutions to DOE mission problems in the 2025–2030 timeframe. On June 28, 2022, a Request for Information (RFI) was issued to gather input from vendors on approaches for DOE and the labs to follow in securing the next generation of supercomputing systems.

ALCF worked along with the other DOE laboratories to craft the evaluation plans and integration of reviews. After responses were received from roughly 50 organizations, ALCF provided written feedback on the responses assigned to Argonne and worked with the other labs to produce a summary. Subsequently, the vendors were invited to participate in a virtual roundtable to discuss priorities for future R&D investment and ways to improve vendor engagement and procurement models. The (non-NDA) roundtable was held on December 12, 2022, with 30-minute Q&A sessions with individual vendors under NDAs held on December 12 and 16, 2022, and January 11, 2023. ALCF helped plan the vendor roundtable and attended all the vendor meetings.

The next step is for the six laboratories to use the RFI responses and feedback in drafting an Advanced Computing Technology Development & Deployment Strategy.

8.3 DOE Program Engagements / Requirements Gathering

To help ensure that the ALCF delivers on its mission of delivering breakthrough science, staff need to closely engage with domain science and keep a close eye on directions of supercomputing technologies. ALCF provides a crucial insight into how production science applications and computer science technologies can move into new machine architectures in the near term and longer term.

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 production resources and future resources alike. Additionally, staff members participate in community and domain activities, including conferences, workshops, reviews, and meetings. In CY 2022, staff participated in more than 191 events. Figure 8.6 breaks down these events by both type and community. Staff members support DOE mission needs by serving on review committees and advisory boards and by participating and organizing DOE and broader community workshops. ALCF staff are regular participants in DOE and NSF workshops and reviews. Staff are engaged in standards committees and boards for both future and current software and hardware technologies.

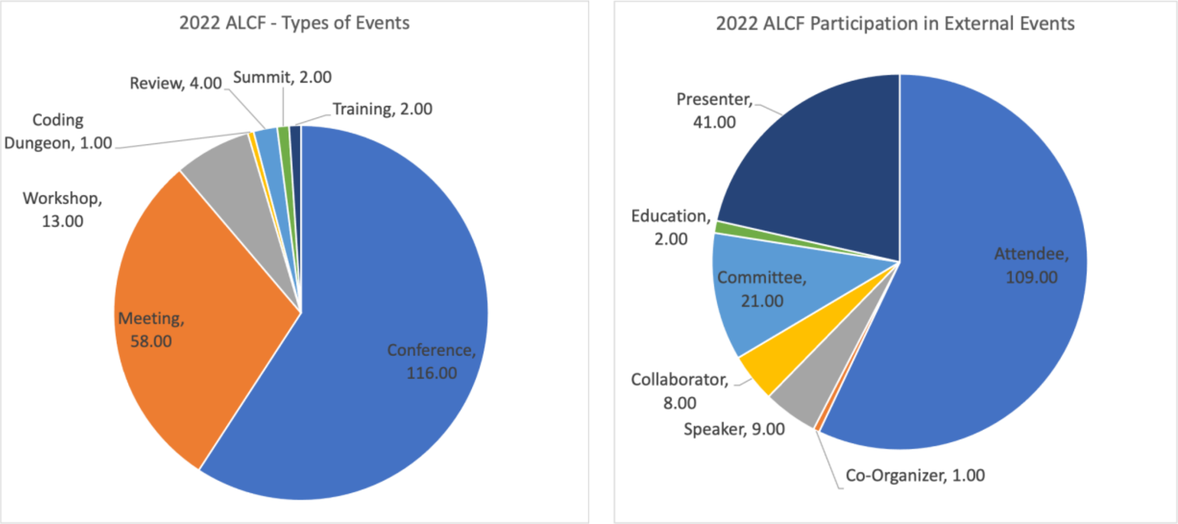

Figure 8.6 Breakdown of some key activities by ALCF in CY 2022. The first pie chart (left) breaks down the 191 events by type, primarily derived by how the event identified itself. The second pie chart (right) breaks down the same events by the role of the staff member.

Not only do these activities maintain expertise of the staff, but they show the respect that ALCF staff have in the community.

8.3.1 Engagement Highlights

Supercomputing 2022

SC is one of the key events in the field of supercomputing and covers every area of the field. Participation is one of the primary opportunities to document and share key knowledge. ALCF has significant participation in the event, as shown in Table 8.3.

Table 8.3 Summary of ALCF Participation in SC22

Program	Total
Tutorials	2
Workshops	14
Papers	2
Posters	2
Birds of a Feather	5
Booth Talks	6

Engagement in Standards and Community Groups

ALCF staff members remain actively involved in numerous HPC standards and community groups that help drive improvements in the usability and efficiency of scientific computing tools, technologies, and applications. Staff activities include contributions to the C++ Standards Committee, Cray User Group, HPC User Forum, Intel eXtreme Performance Users Group (IXPUG), Khronos OpenCL and SYCL Working Groups, MLPerf (HPC, Science, and Storage Working Groups), MPI Forum, NITRD Middleware and Grid Infrastructure Team, oneAPI Community Forum (Steering Committee; Hardware, Language, and Math Special Interest Groups), OpenMP Architecture Review Board, OpenMP Language Committee, Open Scalable File Systems (OpenSFS) Board, and SPEC (Standard Performance Evaluation Corporation) HPG (High Performance Group).

Performance, Portability, and Productivity in HPC (P3HPC) Forum

The Performance, Portability and Productivity in HPC workshop at SC22 was jointly organized by ALCF, LLNL, Intel, and NVIDIA to bring together developers and researchers with an interest in the development of performance-portable applications across current and future high-performance computers. The topic of performance, portability, and productivity focuses on enabling applications and libraries to run across multiple architectures without significant impact on achieved performance and with the goal of maintaining developer productivity. It is important that developers understand and enhance the best practices in this area in order to enable applications to run efficiently across the diverse hardware platforms that exist today.

8.3.2 Summary of Engagements with the Exascale Computing Project

Argonne is a core laboratory of the ECP, and several members of ALCF's leadership team are engaged in the ECP project. Susan Coghlan and David Martin are a part of the ECP leadership team: Coghlan is deputy director of Hardware and Integration (HI), and Martin is co-executive director of the ECP Industry and Agency Council. Haritha Siddabathuni Som is the level-3 lead for Facility Resource Utilization, and Scott Parker is the level-3 lead for Application Integration. Christopher Knight is a level-4 lead for the Aurora Application Integration area. Other leadership team members participate in the various working groups, including Bill Allcock and Jini Ramprakash. 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.

In CY 2022, 25 ALCF staff members attended the Virtual ECP Annual Meeting held May 2–6, 2022, to participate in technical conversations, project discussions, and facility-specific breakouts. In addition, ALCF participated in several planning meetings with ECP and the other computing facilities (NERSC, OLCF) to augment and execute the ECP/Facilities engagement plan and worked with ECP's Training Lead to promote ECP training activities to ALCF users.

ECP-Funded Positions in ALCF

The ALCF's ECP Hardware & Integration effort made great strides in 2022, continuing the team's work in porting and testing ECP applications across many GPUs, including Intel GPUs. There are 37 ALCF staff members funded at various levels to work with ECP Application Development and Software Technology projects. One staff member focuses on training, and Intel's Center of Excellence (COE) for Aurora is staffed with six people. Additionally, ten staff members have been funded to develop and deploy ECP continuous integration (CI) capabilities, support software technologies, and work with others within ECP on containers. Two contractors were hired to develop specific enhancements to the Gitlab-CI platform used for continuous integration. As new ECP project teams were onboarded at JLSE, additional staff members and a contractor were funded to support these project teams. Five staff members were funded to explore the HPE/Cray Shasta software stack. Finally, fifteen staff members were funded to aid ECP applications in porting and testing.

Continuous Integration (CI) Pipeline

In 2022, ALCF continued to support the growth of the CI Pipeline through the ECP-CI project using Gitlab-CI. This resource provides a key tool for projects to perform regular, automated testing on ALCF resources. Through Gitlab-CI, users can enable their CI pipelines on Theta as well as on early hardware available through JLSE. Users ran more than 30,000 jobs within the Gitlab-CI infrastructure with 84% of those jobs completing successfully. In 2023, ALCF will enable Gitlab-CI access to Polaris and open up access more broadly to the user community.

Communication between the ALCF and the ECP Resource Allocation Council

In 2018, the ECP ALCC allocation ended, and the DOE computing facilities switched to the Resource Allocations Council (RAC) to support ECP computing needs. The RAC, 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 sends allocation and usage data in CSV (comma separated values) files to the ECP each day (one for Theta and one for Polaris). The files are uploaded to a Box folder accessible by ECP from where it is downloaded, processed, and merged into the data pipeline that feeds into the ECP User Program dashboard.

Conclusion

The ALCF continues to enable scientific achievements, consistent with DOE's strategic goals of scientific breakthroughs and foundations of science, through projects carried out on ALCF resources. In 2022, researchers participating in projects using ALCF resources published 213 papers in high-quality conferences and journals. ALCF projects have had success in a variety of fields, using many different computational approaches. ALCF projects have been able to reach their scientific goals and successfully use their allocations. Several of the projects and

PIs subsequently received awards or were recognized as achieving significant accomplishments in their fields.

This page intentionally left blank.

Appendix A – Calculations

A.1 Scheduled Availability

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 by more than 4 hours will be counted as an adjacent unscheduled outage, as an unscheduled availability, and as an additional interrupt.

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

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)

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:

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

A.4 System Mean Time to Failure (MTTF)

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

Formula:

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

A.5 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. Jobs that ran during an outage are excluded.

Formula:

$$\mathbf{Utilization} = \left(\frac{\boldsymbol{Node\ Hours\ used\ in\ period}}{\boldsymbol{Node\ Hours\ available\ in\ period}}\right) * \mathbf{100}$$

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.

Table A.1 shows the capability definitions for reportable machine Theta.

Table A.1 Capability Definitions for Theta

Theta				
Capability	**High Capability**	**Range**	**Minimum Nodes**	**Maximum Nodes**
No	No	0% <= x < 20.0%	1	799
Yes	No	20.0% <= x < 60.0%	800	2,399
Yes	Yes	60.0% <= x	2,400	See: A.7 Theta Nodes

Capability also refers to a calculation. The capability calculation is the percentage of node-hours of jobs with the capability attribute versus the total node-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 node-hours of INCITE jobs with the capability attribute versus the total node-hours of all INCITE jobs for a time period.

***Formula*:**

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

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

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.2.

Table A.2 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 node-hours for each day was calculated as the minimum of the node-hours used and the node-hours possible.
- Overall Capability: 20 percent of the reportable nodes.
- High Capability: 60 percent of the reportable nodes.

This page intentionally left blank.

Appendix B – ALCF Director's Discretionary Projects

August 9, 2022–December 31, 2022

Director's Discretionary (DD) Projects on Polaris

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
ACO2RDS	John J. Low	Argonne National Laboratory	Adsorptive CO2 Removal from Dilute Sources	Materials Science	1,250
AI4NMR	Eric Michael Jonas	The University of Chicago (UChicago)	Structure Elucidation for Nuclear Magnetic Resonance via Structured Prediction	Chemistry	5,337
ALCFAITP	Venkatram Vishwanath	Argonne National Laboratory	Argonne AI Training Program	Training	6,000
alcf_training	Yasaman Ghadar	Argonne National Laboratory	ALCF Training	Training	150
AMRdetonations	Venkatramanan Raman	University of Michigan	Adaptive Simulation of Detonations	Engineering	796
APSDataAnalysis	Rafael Vescovi	Argonne National Laboratory	APS Beamline Data Processing and Analysis	Computer Science	1,726
APSDataProcessing	Nicholas Schwarz	Argonne National Laboratory	Advanced Photon Source (APS) Data Processing	Computer Science	1,894
atlas_aesp	Walter Howard Hopkins	Argonne National Laboratory	Simulating and Learning in the ATLAS detector at the Exascale	Physics	2,730
ATPESC2022	Raymond M. Loy	Argonne National Laboratory	Argonne Training Program for Extreme-Scale Computing 2022	Computer Science	1,800
autopology_alcf	Rafael Gomez-Bombarelli	Massachusetts Institute of Technology (MIT)	End to End Classical Force Field Parametrization for Polymer Electrolytes Using Machine Learning	Materials Science	723
BFTrainer	Rajkumar Kettimuthu	Argonne National Laboratory	Rescaling DNN Training Tasks to Fit Dynamically Changing Holes in Supercomputer Schedule	Computer Science	3,034
BPC	Christopher Michael Graziul	The University of Chicago (UChicago)	Optimization of Audio Processing Pipeline for Broadcast Police Communications	Computer Science	4,470
BRAIN	Getnet Dubale Betrie	Argonne National Laboratory	Scalable Brain Simulator for Extreme Computing	Biological Sciences	4,000
candle_aesp	Rick Lyndon Stevens, Thomas Scott Brettin	Argonne National Laboratory	Virtual Drug Response Prediction	Biological Sciences	10,922
catalysis_aesp	David Hamilton Bross	Argonne National Laboratory	Exascale Computational Catalysis	Chemistry	5,461

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
Catalyst	Katherine M. Riley, Christopher James Knight, James Clifton Osborn, Timothy Joe Williams	Argonne National Laboratory	Catalyst	Internal	27,410
cfdml_aesp	Kenneth Edward Jansen	University of Colorado-Boulder	Data Analytics and Machine Learning for Exascale CFD	Engineering	5,461
CFS_UX_TEST	Haritha Siddabathuni Som	Argonne National Laboratory	TESTING CFS	Support	8
climate_downscale	Veerabhadra Rao Kotamarthi	Argonne National Laboratory	Regional Scale Climate Modeling	Earth Science	4,000
co2-ads-mof	Jonathan Rutherford Owens	General Electric Company (GE)	Understanding CO2 Adsorption on Metal-Organic-Frameworks	Chemistry	1,068
connectomics_aesp	Nicola Joy Ferrier, Thomas David Uram	Argonne National Laboratory	Enabling Connectomics at Exascale to Facilitate Discoveries in Neuroscience	Biological Sciences	273
covid-ct	Ravi Kiran Madduri	Argonne National Laboratory	Medical Imaging Domain-Expertise Machine Learning for Interrogation of COVID	Computer Science	2,427
Cray	Torrance Ivan Leggett, Mark Richard Fahey, Susan Marie Coghlan, Timothy Joe Williams, William Edward Allcock	Hewlett Packard Enterprise	Cray Installation	Internal	22,656
CSC249ADCD01	Ian Foster	Argonne National Laboratory	2.2.6.03 ADCD01-CODAR	Computer Science	1,716
CSC249ADCD02	Susan Marie Mniszewski, Timothy C. Germann	Los Alamos National Laboratory (LANL)	2.2.6.04 ADCD02-COPA: Co-Design Center for Particle Applications	Physics	1,716
CSC249ADCD04	Tzanio Valentinov Kolev, Misun Min, Paul Frederick Fischer	Lawrence Livermore National Laboratory (LLNL)	2.2.6.06 CEED: Center for Efficient Exascale Discretizations	Computer Science	6,716
CSC249ADCD05	Mahantesh Halappanavar	Pacific Northwest National Laboratory (PNNL)	2.2.6.07 ADCD05-ExaGraph	Computer Science	10,716
CSC249ADCD08	Francis Joseph Alexander	Brookhaven National Laboratory (BNL)	2.2.6.08 ADCD08-ExaLearn	Physics	1,716
CSC249ADCD09	John Bell	Lawrence Berkeley National Laboratory (LBNL)	2.2.6.05 ADCD03-AMREX: Block-Structured AMR Co-Design Center	Mathematics	110
CSC249ADCD502	Kenneth John Roche	Pacific Northwest National Laboratory (PNNL)	2.2.6.02 ADCD502 Application Assessment	Computer Science	1,716

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CSC249ADCD504	Jeanine Cook, Shirley Victoria Moore	Lawrence Livermore National Laboratory (LLNL)	2.2.6.01 ADCD504-Proxy Applications	Computer Science	1,716
CSC249ADOA01	Rick Lyndon Stevens, Thomas Scott Brettin	Argonne National Laboratory	2.2.4.03 ADOA01 CANDLE: Exascale Deep Learning Enabled Precision Medicine for Cancer	Biological Sciences	13,716
CSC249ADSE03	Andreas Samuel Kronfeld, Norman Howard Christ, Paul Mackenzie	Fermi National Accelerator Laboratory (Fermilab)	2.2.1.01 ADSE03-LatticeQCD: Exascale Lattice Gauge Theory Opportunities/Reqmts for Nuclear & High Energy Physics	Physics	1,716
CSC249ADSE04	Danny Perez	Los Alamos National Laboratory (LANL)	2.2.1.04 ADSE04-EXAALT - Molecular Dynamics at the Exascale	Nuclear Energy	1,716
CSC249ADSE05	David Paul Trebotich	Lawrence Berkeley National Laboratory (LBNL)	2.2.3.04 ADSE05-Subsurface	Earth Science	1,716
CSC249ADSE06	Jean-Luc Yves Vay	Lawrence Berkeley National Laboratory (LBNL)	2.2.2.06 ADSE06-WarpX	Physics	1,716
CSC249ADSE08	Steven Hamilton, Paul Kollath Romano	Oak Ridge National Laboratory (ORNL)	2.2.2.03 ADSE08 ExaSMR	Nuclear Energy	1,716
CSC249ADSE09	Paul Richard Charles Kent, Anouar Benali	Oak Ridge National Laboratory (ORNL)	2.2.1.06 QMCPACK: Predictive and Improvable Quantum-mechanics Based Simulations	Materials Science	1,716
CSC249ADSE11	Theresa Windus	University of Washington	2.2.1.02 ADSE11-NWChemEx: Tackling Chemical, Materials, & Biomolecular Challenges in Exascale	Chemistry	1,716
CSC249ADSE12	Amitava Bhattacharjee	Princeton Plasma Physics Laboratory (PPPL)	2.2.2.05 ADSE12 WDMAPP	Computer Science	1,716
CSC249ADSE14	Jacqueline Chen	Sandia National Laboratories, California	2.2.2.02 ADSE14-Combustion-Pele: Transforming Combustion Science & Technology with Exascale Simulations	Engineering	1,716
CSC249ADSE15	Mark Alan Taylor	Sandia National Laboratories, New Mexico	2.2.3.05 ADSE15-E3SM-MMF	Earth Science	110
CSC249ADSE16	Mark S. Gordon	Ames Laboratory	2.2.1.03 ADSE16-GAMESS	Chemistry	1,716
CSC249ADSE18	Daniel Kasen	Lawrence Berkeley National Laboratory (LBNL)	2.2.3.01 ADSE18 Exastar	Physics	1,716

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CSC249ADSE22	Christopher Stephen Oehmen, Andres Marquez, Zhenyu Huang	Pacific Northwest National Laboratory (PNNL)	2.2.4.02 ADSE22-ExaSGD	Energy Technologies	1,716
CSC249ADSE23	Jordan Michael Musser	National Energy Technology Laboratory (NETL)	2.2.2.04: MFIX-Exa: Perf Prediction of Multiphase Energy Conversion Device	Energy Technologies	1,000
CSC249ADTR01	Daniel Edward Laney	Lawrence Livermore National Laboratory (LLNL)	2.3.5.10 ADTR01-ExaWorks	Computer Science	1,716
CSC249ADTR02	Osni Marques	Oak Ridge National Laboratory (ORNL)	2.4.6.02 ADTR02- Productivity	Computer Science	1,716
CSC250STDA05	Kenneth Dean Moreland	Oak Ridge National Laboratory (ORNL)	2.3.4.13 STDA05-ECP/VTK-m	Computer Science	1,716
CSC250STDM10	Surendra Byna, Venkatram Vishwanath	Lawrence Berkeley National Laboratory (LBNL)	2.3.4.15 ExaIO - Delivering Efficient Parallel I/O on Exascale Computing Systems with HDF5 and Unify	Computer Science	10,716
CSC250STDM11	Scott Klasky, Norbert Podhorszki	Oak Ridge National Laboratory (ORNL)	2.3.4.09 STDM11-ADIOS Framework for Scientific Data on Exascale Systems	Computer Science	1,716
CSC250STDM12	Robert B. Ross, Robert J. Latham	Argonne National Laboratory	2.3.4.10 STDM12-DataLib: Data Libraries and Services Enabling Exascale Science	Computer Science	1,716
CSC250STDM14	Franck Cappello	Argonne National Laboratory	2.3.4.14 STDM14 - VeloC-SZ: Very Low Overhead Transparent Multilevel Checkpoint/Restart/SZ: Fast, Effective, Parallel Error-bounded Exascale Loss....	Computer Science	1,716
CSC250STDM16	James Paul Ahrens, Terece Louise Turton	Lawrence Livermore National Laboratory (LLNL)	2.3.4.16 STDM16-ALPINE/ZFP	Computer Science	1,716
CSC250STDT10	Jeffrey S. Vetter	Oak Ridge National Laboratory (ORNL)	2.3.2.10 STDT10 PPROTEAS-TUNE	Computer Science	1,716
CSC250STDT11	Sunita Chandrasekaran, Dossay Oryspayev	Stony Brook University	2.3.2.11 SOLLVE: Scaling OpenMP with LLVm for Exascale	Computer Science	1,716
CSC250STDV01	Charles Vernon Atkins	Kitware Inc.	2.3.4.01 STDV01-Data and Visualization Software Development Kit	Computer Science	1,716
CSC250STML12	Carol Woodward	Lawrence Livermore National Laboratory (LLNL)	2.3.3.12 Enabling Exascale Simulations with SUNDIALS and hypre	Mathematics	2,000
CSC250STML13	Hartwig Andreas Anzt	The University of Tennessee at Knoxville	2.3.3.13 STML13 - CLOVER	Computer Science	1,000

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CSC250STML15	Siva Rajamanickam	Sandia National Laboratories, New Mexico	2.3.3.15 STML-Sake	Computer Science	2,000
CSC250STMS05	Ulrike Meier Yang, Satish Balay	Argonne National Laboratory	2.3.3.01 STMS05-Extreme-scale Scientific xSDK for ECP	Mathematics	1,716
CSC250STMS07	Todd S. Munson, Hong Zhang, Richard Tran Mills, Satish Balay	Argonne National Laboratory	2.3.3.06 STMS07-PETSc/TAO for Exascale	Mathematics	1,716
CSC250STMS08	Xiaoye Sherry Li	Lawrence Berkeley National Laboratory (LBNL)	2.3.3.07 STMS08 STRUMPACK/SuperLU/FFTX: Factorization Based Sparse Solvers and Preconditioners for Exascale	Mathematics	100
CSC250STNS01	Michael Lang, Terece Louise Turton	Los Alamos National Laboratory (LANL)	2.3.6.01 - STNS01 -LANL ATDM ST Projects	Computer Science	1,716
CSC250STPM01	Sameer Suresh Shende	University of Oregon	2.3.1.01 Programming Models & Runtimes Software Development Kit	Computer Science	1,000
CSC250STPM09	Yanfei Guo	Argonne National Laboratory	2.3.1.07 STPM09-Exascale MPI	Computer Science	1,716
CSC250STPM11	George Bosilca, Earl Luther Carr, Jack Dongarra, Thomas Herault	The University of Tennessee at Knoxville	2.3.1.09 STPM11 ParSEC: Distributed Tasking	Computer Science	1,716
CSC250STPM16	Latchesar Alexandrov Lonkov	Los Alamos National Laboratory (LANL)	2.3.1.16 SICM: Simplified Interface to Complex Memory	Computer Science	1,000
CSC250STPM17	Paul Hamilton Hargrove, Erich Strohmaier	Lawrence Berkeley National Laboratory (LBNL)	2.3.1.14 STPM17-UPC++ & GASNet	Computer Science	2,216
CSC250STPM18	Christian Trott	Sandia National Laboratories, California	2.3.1.18 RAJA/Kokkos	Computer Science	1,716
CSC250STPR19	Peter Hugh Beckman	Argonne National Laboratory	2.3.1.19 STPR19 Argo: Argo/Power Steering	Computer Science	1,716
CSC250STPR27	David Edward Bernholdt	Oak Ridge National Laboratory (ORNL)	2.3.1.17 STPR27-OMPI-X: Open MPI for Exascale	Materials Science	1,716
CSC250STTO09	Hartwig Andreas Anzt, Anthony Danalis, Earl Luther Carr, Heike Jagode	The University of Tennessee at Knoxville	2.3.2.06 STTO09 EXAPAPI	Computer Science	1,716
CSC250STTO11	John Michael Mellor-Crummey	Rice University	2.3.2.08 STTO11 HPCToolkit	Computer Science	1,716
CSC251HIHE05	Scott Dov Pakin, Simon David Hammond	Los Alamos National Laboratory (LANL)	2.4.2.01 HIHE05-Analytical Modeling - Hardware Evaluation Working Groups	Computer Science	1,716

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CSC251HISD01	Ryan Charles Prout	Los Alamos National Laboratory (LANL)	2.4.4.01 HISD01-Software Integration	Computer Science	1,716
CSCSTDT12345	Patrick McCormick	Los Alamos National Laboratory (LANL)	2.3.2.12 Flang: Open-Source Fortran Front End for the LLVM Infrastructure	Computer Science	1,716
darkskyml_aesp	Salman Habib	Argonne National Laboratory	Dark Sky Mining	Physics	273
datascience	Venkatram Vishwanath	Argonne National Laboratory	ALCF Data Science and Workflows Allocation	Internal	27,410
determined_eval	Venkatram Vishwanath	Argonne National Laboratory	Evaluation of Determined.AI HPO on Polaris	Computer Science	500
DLHMC	Sam Alfred Foreman	Argonne National Laboratory	Deep Learning HMC	Physics	2,144
DNS3D	Ramesh Balakrishnan	Argonne National Laboratory	Direct Numerical Simulation of Three-Dimensional Turbulence	Engineering	292
DNSVCMHD	Keith Daniel Brauss	Francis Marion University	DNS Simulations of Velocity-Current Magnetohydrodynamic Equations	Mathematics	1,146
Drag-Reduction	Paul Fischer	University of Illinois at Urbana-Champaign	DNS of Drag Reduction	Engineering	2,363
EE-ECP	Xingfu Wu, Valerie Taylor	Argonne National Laboratory	Energy Efficient Tradeoff among Execution Time and Power of ECP Applications	Computer Science	750
ESGF2	Ian Foster	Argonne National Laboratory	ESGF2	Earth Science	257
FLUPS	Gilles Poncelet	Université Catholique de Louvain	3D Distributed Fourier-based Poisson Solver	Engineering	1,181
fusiondl_aesp	William Tang	Princeton University	Accelerated Deep Learning Discovery in Fusion Energy Science	Fusion Energy	273
GPUBenchDFT	Ganesh Sivaraman	Argonne National Laboratory	Benchmark of GPU Based Real Space and Plane Wave DFT Codes	Chemistry	2,363
gpu_hack	Yasaman Ghadar, Raymond M. Loy	Argonne National Laboratory	GPU Hackathon	Training	522
GRACE	Sayan Ghosh	Pacific Northwest National Laboratory (PNNL)	Graph Analytics Codesign on GPUs	Computer Science	1,021
HACC_aesp	Katrin Heitmann	Argonne National Laboratory	Extreme-Scale Cosmological Hydrodynamics	Physics	273
hp-ptycho	Tekin Bicer	Argonne National Laboratory	High Performance 3D Ptychographic Reconstruction and Image Enhancement	Materials Science	1,321

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
hydrosm	Jeremy A. Feinstein	Argonne National Laboratory	Improving the Predictability of Hydrological Systems with AI	Earth Science	343
hyper_bound_layer	Carlo Scalo	Purdue University	Passive Control of Hypersonic Boundary Layers via Wall Treatments	Engineering	1,149
IBM-GSS	Venkatram Vishwanath	Argonne National Laboratory	IBM GeoSpatial Software System	Earth Science	1,308
Intel	Kalyan Kumaran, Scott Parker, Timothy Joe Williams, Venkatram Vishwanath	Argonne National Laboratory	Intel Employees in Support of Theta	Internal	711
LASSCF_gpudev	Christopher James Knight	Argonne National Laboratory	GPU Development of LASSCF	Chemistry	597
LatticeQCD_aesp	Paul Mackenzie, Norman Howard Christ	Fermi National Accelerator Laboratory (Fermilab)	Lattice Quantum Chromodynamics Calculations for Particle and Nuclear Physics	Physics	273
LQCDdev	James Clifton Osborn	Argonne National Laboratory	Lattice QCD Development	Physics	248
lqcdml_aesp	William Detmold	Massachusetts Institute of Technology (MIT)	Machine Learning for Lattice Quantum Chromodynamics	Physics	7,509
Maintenance	William Edward Allcock, John Francis O'Connell, John Patrick Reddy, Ryan Milner, Torrance Ivan Leggett	Argonne National Laboratory	LCF Operations System Maintenance	Internal	5,336
matml_aesp	Noa Marom	Carnegie Mellon University	Many-Body Perturbation Theory Meets Machine Learning to Discover Singlet Fission Materials	Materials Science	6,553
MFIX-Exa	William David Fullmer	National Energy Technology Laboratory (NETL)	MFIX-Exa: Performance Prediction of Multiphase Energy Conversion Device	Energy Technologies	1,402
MLP4THERMO	Cem Sevik	Eskisehir Technical University	Machine Learning Potentials for Thermal Properties of Two-Dimensional Materials	Materials Science	2,043
MOAB_App	Vijay Subramaniam Mahadevan	Argonne National Laboratory	MOAB Algorithmic Performance Portability	Mathematics	2,243
MultiActiveAI	Dario Dematties	Northwestern Argonne Institute of Science and Engineering (NAISE)	Multimodal Intelligence for Federated Edge Computing Simulations	Computer Science	171
multiphysics_aesp	Amanda Randles	Duke University	Extreme-scale In Situ Visualization and Analysis of Fluid-Structure-Interaction Simulations	Engineering	7,372

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
MVAPICH2	Dhabaleswar Kumar Panda	The Ohio State University	Optimizing and Tuning MVAPICH2-GDR Library and Study Its Impact on HPC and AI Applications	Computer Science	2,105
NAMD_aesp	Benoit Roux, James Christopher Phillips	The University of Chicago (UChicago)	Free Energy Landscapes of Membrane Transport Proteins	Biological Sciences	5,461
NAQMC_RMD_aesp	Aiichiro Nakano	University of Southern California (USC)	Metascalable Layered Materials Genome	Materials Science	273
nekrs-scaling	Pinaki Pal	Argonne National Laboratory	NekRS Scalability Studies for Gas Turbine Film Cooling High-Fidelity Simulations	Energy Technologies	1,181
NWChemEx_aesp	Theresa Windus, Alvaro Vazquez Mayagoitia	Pacific Northwest National Laboratory (PNNL)	NWChemEx: Tackling Chemical, Materials & Biochemical Challenges in the Exascale Era	Chemistry	6,826
OmniverseEval	Joseph A. Insley	Argonne National Laboratory	NVIDIA Omniverse Evaluation	Computer Science	2,200
Operations	William Edward Allcock	Argonne National Laboratory	Systems Administration Tasks	Internal	27,410
PARTURB3D	Ramesh Balakrishnan	Argonne National Laboratory	Simulating Turbulent Particulate Flows Inside Enclosures	Engineering	468
Performance	Scott Parker, Raymond M. Loy	Argonne National Laboratory	Performance	Internal	27,410
PHASTA_aesp	Kenneth Edward Jansen	University of Colorado-Boulder	Extreme Scale Unstructured Adaptive CFD: From Multiphase Flow to Aerodynamic Flow Control	Engineering	5,461
Polaris	Torrance Ivan Leggett	Hewlett Packard Enterprise	Polaris project for installation and related work for the vendors	Internal	205
QMCPACK_aesp	Anouar Benali	Argonne National Laboratory	Extending Moore’s Law Computing with Quantum Monte Carlo	Materials Science	7,372
QTensor	Yuri Alexeev	Argonne National Laboratory	Quantum Circuit Simulations	Computer Science	2,190
QuantumDS	Alvaro Vazquez Mayagoitia	Argonne National Laboratory	Quantum Mechanics and Data Science	Chemistry	594
radix-io	Philip Hutchinson Carns	Argonne National Laboratory	System Software to Enable Data-Intensive Science	Computer Science	124
RAPINS	Eliu Antonio Huerta Escudero	Argonne National Laboratory	Reproducible and Accelerated Physics-inspired Neural Networks	Physics	214
remote_offloading	Jose Manuel Monsalve Diaz	Argonne National Laboratory	Exploring Collective Operations with Remote Offloading	Computer Science	1,181

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
RL-fold	Arvind Ramanathan	Argonne National Laboratory	Targeting Intrinsically Disordered Proteins Using Artificial Intelligence Driven Molecular Simulations	Biological Sciences	58,657
safcomb	Marcus Steven Day	National Renewable Energy Laboratory (NREL)	Turbulent DNS of SAFs in an Aero Combustor	Chemistry	2,309
sbi-fair	Pete Beckman, Kamil Antoni Iskra	Argonne National Laboratory	FAIR Surrogate Benchmarks Supporting AI and Simulation Research	Computer Science	1,633
scalablepinns	Paris Perdikaris	University of Pennsylvania (UPenn)	Scalable PINNs	Engineering	1,720
SCPlasma	Ranganathan Gopalakrishnan	University of Memphis	Thermodynamics and Transport Models of Strongly Coupled Dusty Plasmas	Physics	343
SCREAM_Calib	Jiali Wang	Argonne National Laboratory	Towards Neighborhood Scale Climate Simulations using AI and Accelerated GPUs	Earth Science	1,068
SDL_Workshop	Yasaman Ghadar	Argonne National Laboratory	ALCF Simulation, Data, and Learning Workshop	Training	12,288
SEEr-planning	Zhiling Lan	Illinois Institute of Technology (IIT)	Performance and Power Tradeoff Analysis of AI-Enabled Science on CPU-GPU System	Computer Science	1,894
SEEr-Polaris	Zhiling Lan	Illinois Institute of Technology (IIT)	AI-enabled Benchmarking on Polaris	Computer Science	2,243
SENSEI	Silvio Humberto Rafael Rizzi, Joseph A. Insley, Nicola Joy Ferrier, Venkatram Vishwanath	Argonne National Laboratory	Scalable Analysis Methods and In Situ Infrastructure for Extreme Scale Knowledge Discovery	Computer Science	268
SolarWindowsADSP	Jacqueline Manina Cole	University of Cambridge	Data-Driven Molecular Engineering of Solar-powered Windows	Materials Science	1,988
STlearn	Shinjae Yoo	Brookhaven National Laboratory (BNL)	Spatiotemporal Learning for Human Neuroscience	Biological Sciences	2,293
Substrate-transport	Wonpil Im	Lehigh University	Characterizing Energy Landscape of Substrate Translocation in Bacterial Membrane	Biological Sciences	248
SuperBERT	Ian Foster	Argonne National Laboratory	Training of Language Models on Large Quantities of Scientific Text	Computer Science	7,959
swift-t-polaris	Justin Michael Wozniak	Argonne National Laboratory	Swift/T on Polaris	Computer Science	138

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
SYCLSupport	Kevin Harms	Argonne National Laboratory	SYCL Support on ALCF Systems	Computer Science	434
ThroughFocal_DD	Jonathan Tyler Schwartz	University of Michigan	Aberration Corrected Through Focal Electron Tomography Through Focal_DD	Materials Science	879
training_polaris	Paige Carolyn Kinsley, Yasaman Ghadar	Argonne National Laboratory	Training Polaris	Training	118
transformer_eval	Rick Lyndon Stevens	Argonne National Laboratory	Transformers on AI Accelerators	Computer Science	487
UINTAH_aesp	Martin Berzins, John Andrew Schmidt	The University of Utah	Design and Evaluation of High-efficiency Boilers for Energy Production Using a Hierarchical V/UQ Approach	Chemistry	273
VeloC	Bogdan Florin Nicolae	Argonne National Laboratory	VeloC: Very Low Overhead Checkpointing System	Computer Science	4,000
visualization	Joseph A. Insley, Michael E. Papka	Argonne National Laboratory	Visualization and Analysis Research and Development for ALCF	Internal	27,410
wall_turb_dd	Ramesh Balakrishnan	Argonne National Laboratory	Wall Resolved Simulations of Canonical Wall Bounded Flows	Engineering	268
wereszczynski	Jeffery Michael Wereszczynski	Illinois Institute of Technology (IIT)	MD Simulations of Chromatin Modification and Gene Regulation Mechanism	Biological Sciences	2,243
WMLES	Zhi Jian WANG	Kansas State University	Wall Modeled Large Eddy Simulation for Turbomachinery Applications	Engineering	2,243
XGC_aesp	Choongseok Chang	Princeton Plasma Physics Laboratory (PPPL)	High Fidelity Simulation of Fusion Reactor Boundary Plasmas	Fusion Energy	5,461
				Total DD	543,571

January 1, 2022–December 31, 2022
Director's Discretionary (DD) Projects on Theta

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
2D-magnet	Can Ataca	University of Maryland, Baltimore County	Correlated Two Dimensional Magnets at Chemical Accuracy	Materials Science	32,000
3DChromatin	Jie Liang	University of Illinois at Chicago	Large Ensemble Model of Single-Cell 3D Genome Structures	Biological Sciences	102,191
ACO2RDS	John J. Low	Argonne National Laboratory	Adsorptive CO2 Removal from Dilute Sources	Materials Science	117,512
ALCFAITP	Venkatram Vishwanath	Argonne National Laboratory	Argonne AI Training Program	Training	1,500
Allinea	Raymond M. Loy, Kalyan Kumaran	Argonne National Laboratory	Improved Debugging Memory Usage for BG/Q	Internal	781
AnMod	Alon Grinberg Dana	Technion - Israel Institute of Technology	Accurate Partition Function Calculation Considering Anharmonic Modes of Complex Chemical Systems	Chemistry	32,000
APSDataAnalysis	Rafael Vescovi	Argonne National Laboratory	APS Beamline Data Processing and Analysis	Computer Science	25,218
Aramco-PreChamber	Joohan Kim	Argonne National Laboratory	High-Fidelity LES of Turbulent Jet Combustion	Engineering	8,774
arfc-msr-ahtr	Madicken Munk	University of Illinois at Urbana-Champaign	Modeling of Molten Salt Reactor Design, Optimization, and Transient Behavior	Nuclear Energy	50,000
atlas_aesp	Walter Howard Hopkins	Argonne National Laboratory	Simulating and Learning in the ATLAS Detector at the Exascale	Physics	19,500
ATPESC2022	Raymond M. Loy	Argonne National Laboratory	Argonne Training Program for Extreme-Scale Computing 2022	Computer Science	5,000
ATPESC_Instructors	Raymond M. Loy	Argonne National Laboratory	Argonne Training Program on Extreme-Scale Computing for ALL Instructors	Training	500
autopology_alcf	Rafael Gomez-Bombarelli	Massachusetts Institute of Technology (MIT)	End to End Classical Force Field Parametrization for Polymer Electrolytes Using Machine Learning	Materials Science	2,000
AXMAS-Flows	Justin Michael Wozniak	Argonne National Laboratory	AXMAS Scalable Workflows	Physics	2,000
BIP167	Philip Kurian	Howard University	Computing Superradiance and van der Waals Many-body Dispersion Effects for Biomacromolecules	Physics	21,470

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
BLawB	Lucian Ivan	Canadian Nuclear Laboratories	Application of Maximum-Entropy Moment Methods to Turbulent and Multiphase Flow Prediction: Software Package Preparation	Engineering	64,000
bloodflow_dd	Jifu Tan	Northern Illinois University (NIU)	Multiphysics Modeling of Biological Flow with Cell Suspensions	Engineering	24,836
BRAIN	Getnet Dubale Betrie	Argonne National Laboratory	Scalable Brain Simulator for Extreme Computing	Biological Sciences	10,598
BS-SOLCTRA	Esteban Meneses	Costa Rica National High Technology Center	Plasma Physics Simulations for SCR-1 Stellarator	Physics	10,522
CAIDS	Julio Cesar Mendez Carvajal	North Carolina State University (NCSU)	Consistent Averaging Procedure for Solving the Fundamental Equations of Fluid Dynamics	Engineering	4,225
candle_aesp	Rick Lyndon Stevens, Thomas Scott Brettin	Argonne National Laboratory	Virtual Drug Response Prediction	Biological Sciences	2,000
Carbon_composites	Hendrik Heinz	University of Colorado-Boulder	Designing Functional Nanostructures and Carbon-Based Composite Materials	Materials Science	18,908
catalysis_aesp	David Hamilton Bross	Argonne National Laboratory	Exascale Computational Catalysis	Chemistry	12,500
Catalyst	Katherine M. Riley, Christopher James Knight, James Clifton Osborn, Timothy Joe Williams	Argonne National Laboratory	Catalyst	Internal	4,000
CatalystAI	Hieu Anh Doan	Argonne National Laboratory	Accelerated Discovery of Multimetallic Alloy Catalysts for CO2 Conversion via Data-driven Artificial Intelligence	Materials Science	62,500
Cellulose-Simulation	Dewei Qi	Western Michigan University	Molecular Dynamics Simulation of Nano-cellulose	Engineering	70,371
cfdml_aesp	Kenneth Edward Jansen	University of Colorado-Boulder	Data Analytics and Machine Learning for Exascale CFD	Engineering	8,000
CFS_UX_TEST	Haritha Siddabathuni Som	Argonne National Laboratory	TESTING CFS	Support	0
CharmRTS	Laxmikant Kale, Abhinav Bhatele, Juan Jose Galvez-Garcia	University of Illinois at Urbana-Champaign	Charm++ and Its Applications	Computer Science	4,000

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
climate_severe	Vittorio Angelo Gensini	Northern Illinois University (NIU)	Anticipating Severe Weather Events via Dynamical Downscaling	Earth Science	25,835
Climate_Water	Jiali Wang	Argonne National Laboratory	Linking Climate to Water: Implementing a 4km Regional Climate Model with Hydrologic Model Coupling (WRF-Hydro) using Argonne's (etc.)	Earth Science	32,000
Clouds	Ian Foster	Argonne National Laboratory	Unsupervised Analysis of Satellite Cloud Imagery	Earth Science	43,151
CobaltDevel	Paul Michael Rich, William Edward Allcock	Argonne National Laboratory	Cobalt Development	Internal	5,574
COMPASS-GLM	William James Pringle	Argonne National Laboratory	Coastal Observations, Mechanisms, and Predictions Across Systems and Scales – Great Lakes Modeling (COMPASS-GLM)	Earth Science	32,000
Comp_Perf_Workshop	Raymond M. Loy, Yasaman Ghadar	Argonne National Laboratory	ALCF Computational Performance Workshop	Training	100,000
CONUS-Carbon	Jinxun Liu	U.S. Geological Survey (USGS)	Terrestrial Ecosystem Carbon Cycle of the Conterminous U.S.	Earth Science	15,631
CorrVSVB	Graham Donald Fletcher	Argonne National Laboratory	Estimating Correlation Energies Using VSVB	Chemistry	15,625
covid-ct	Ravi Kiran Madduri	Argonne National Laboratory	Medical Imaging Domain-Expertise Machine Learning for Interrogation of COVID	Computer Science	3,363
CPOL	Scott Matthew Collis	Argonne National Laboratory	Learning the Physics of Precipitation from Radar Measurements	Earth Science	32,000
Cray	Torrance Ivan Leggett, Mark Richard Fahey, Susan Marie Coghlan, Timothy Joe Williams, William Edward Allcock	Hewlett Packard Enterprise	Cray Installation	Internal	1,000
cray-hpo	Michael Adnan Salim	Argonne National Laboratory	Scaling Studies of Cray AI Hyperparameter Optimization	Computer Science	200
CSC249ADCD01	Ian Foster	Argonne National Laboratory	2.2.6.03 ADCD01-CODAR	Computer Science	4,000
CSC249ADCD02	Susan Marie Mniszewski, Timothy C. Germann	Los Alamos National Laboratory (LANL)	2.2.6.04 ADCD02-COPA: Co-Design Center for Particle Applications	Physics	4,000

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CSC249ADCD04	Tzanio Valentinov Kolev, Misun Min, Paul Frederick Fischer	Lawrence Livermore National Laboratory (LLNL)	2.2.6.06 CEED: Center for Efficient Exascale Discretizations	Computer Science	11,784
CSC249ADCD05	Mahantesh Halappanavar	Pacific Northwest National Laboratory (PNNL)	2.2.6.07 ADCD05-ExaGraph	Computer Science	14,000
CSC249ADCD08	Francis Joseph Alexander	Brookhaven National Laboratory (BNL)	2.2.6.08 ADCD08-ExaLearn	Physics	41,000
CSC249ADCD502	Kenneth John Roche	Pacific Northwest National Laboratory (PNNL)	2.2.6.02 ADCD502 Application Assessment	Computer Science	4,000
CSC249ADCD504	Jeanine Cook, Shirley Victoria Moore	Lawrence Livermore National Laboratory (LLNL)	2.2.6.01 ADCD504-Proxy Applications	Computer Science	400
CSC249ADOA01	Rick Lyndon Stevens, Thomas Scott Brettin	Argonne National Laboratory	2.2.4.03 ADOA01 CANDLE: Exascale Deep Learning Enabled Precision Medicine for Cancer	Biological Sciences	4,500
CSC249ADSE03	Andreas Samuel Kronfeld, Norman Howard Christ, Paul Mackenzie	Fermi National Accelerator Laboratory (Fermilab)	2.2.1.01 ADSE03-LatticeQCD: Exascale Lattice Gauge Theory Opportunities/Reqmts for Nuclear & High Energy Physics	Physics	4,000
CSC249ADSE04	Danny Perez	Los Alamos National Laboratory (LANL)	2.2.1.04 ADSE04-EXAALT - Molecular Dynamics at the Exascale	Nuclear Energy	400
CSC249ADSE05	David Paul Trebotich	Lawrence Berkeley National Laboratory (LBNL)	2.2.3.04 ADSE05-Subsurface	Earth Science	400
CSC249ADSE06	Jean-Luc Yves Vay	Lawrence Berkeley National Laboratory (LBNL)	2.2.2.06 ADSE06-WarpX	Physics	400
CSC249ADSE08	Steven Hamilton, Paul Kollath Romano	Oak Ridge National Laboratory (ORNL)	2.2.2.03 ADSE08 ExaSMR	Nuclear Energy	400
CSC249ADSE09	Paul Richard Charles Kent, Anouar Benali	Oak Ridge National Laboratory (ORNL)	2.2.1.06 QMCPACK: Predictive and Improvable Quantum-mechanics Based Simulations	Materials Science	4,000
CSC249ADSE11	Theresa Windus	University of Washington	2.2.1.02 ADSE11-NWChemEx: Tackling Chemical, Materials, & Biomolecular Challenges in Exascale	Chemistry	400
CSC249ADSE12	Amitava Bhattacharjee	Princeton Plasma Physics Laboratory (PPPL)	2.2.2.05 ADSE12 WDMAPP	Computer Science	28,000

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CSC249ADSE14	Jacqueline Chen	Sandia National Laboratories, California	2.2.2.02 ADSE14-Combustion-Pele: Transforming Combustion Science & Technology with Exascale Simulations	Engineering	400
CSC249ADSE16	Mark S. Gordon	Ames Laboratory	2.2.1.03 ADSE16-GAMESS	Chemistry	35,500
CSC249ADSE18	Daniel Kasen	Lawrence Berkeley National Laboratory (LBNL)	2.2.3.01 ADSE18 Exastar	Physics	400
CSC249ADSE22	Christopher Stephen Oehmen, Andres Marquez, Zhenyu Huang	Pacific Northwest National Laboratory (PNNL)	2.2.4.02 ADSE22-ExaSGD	Energy Technologies	400
CSC249ADTR01	Daniel Edward Laney	Lawrence Livermore National Laboratory (LLNL)	2.3.5.10 ADTR01-ExaWorks	Computer Science	400
CSC249ADTR02	Osni Marques	Oak Ridge National Laboratory (ORNL)	2.4.6.02 ADTR02-Productivity	Computer Science	400
CSC250STDA05	Kenneth Dean Moreland	Oak Ridge National Laboratory (ORNL)	2.3.4.13 STDA05-ECP/VTK-m	Computer Science	400
CSC250STDM10	Surendra Byna, Venkatram Vishwanath	Lawrence Berkeley National Laboratory (LBNL)	2.3.4.15 ExaIO - Delivering Efficient Parallel I/O on Exascale Computing Systems with HDF5 and Unify	Computer Science	14,500
CSC250STDM11	Scott Klasky, Norbert Podhorszki	Oak Ridge National Laboratory (ORNL)	2.3.4.09 STDM11-ADIOS Framework for Scientific Data on Exascale Systems	Computer Science	4,000
CSC250STDM12	Robert B. Ross, Robert J. Latham	Argonne National Laboratory	2.3.4.10 STDM12-DataLib: Data Libraries and Services Enabling Exascale Science	Computer Science	4,000
CSC250STDM14	Franck Cappello	Argonne National Laboratory	2.3.4.14 STDM14 - VeloC-SZ: Very Low Overhead Transparent Multilevel Checkpoint/Restart/SZ: Fast, Effective, Parallel Error-bounded Exascale Loss	Computer Science	4,000
CSC250STDM16	James Paul Ahrens, Terece Louise Turton	Lawrence Livermore National Laboratory (LLNL)	2.3.4.16 STDM16-ALPINE/ZFP	Computer Science	400
CSC250STDT10	Jeffrey S. Vetter	Oak Ridge National Laboratory (ORNL)	2.3.2.10 STDT10 PPROTEAS-TUNE	Computer Science	400
CSC250STDT11	Sunita Chandrasekaran, Dossay Oryspayev	Stony Brook University	2.3.2.11 SOLLVE: Scaling OpenMP with LLVm for Exascale	Computer Science	400
CSC250STDV01	Charles Vernon Atkins	Kitware Inc.	2.3.4.01 STDV01-Data and Visualization Software Development Kit	Computer Science	14,000

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
CSC250STML15	Siva Rajamanickam	Sandia National Laboratories, New Mexico	2.3.3.15 STML-Sake	Computer Science	1,000
CSC250STMS05	Ulrike Meier Yang, Satish Balay	Argonne National Laboratory	2.3.3.01 STMS05-Extreme-scale Scientific xSDK for ECP	Mathematics	400
CSC250STMS07	Todd S. Munson, Hong Zhang, Richard Tran Mills, Satish Balay	Argonne National Laboratory	2.3.3.06 STMS07-PETSc/TAO for Exascale	Mathematics	4,000
CSC250STNS01	Michael Lang, Terece Louise Turton	Los Alamos National Laboratory (LANL)	2.3.6.01 - STNS01 -LANL ATDM ST Projects	Computer Science	400
CSC250STPM09	Yanfei Guo	Argonne National Laboratory	2.3.1.07 STPM09-Exascale MPI	Computer Science	400
CSC250STPM11	George Bosilca, Earl Luther Carr, Jack Dongarra, Thomas Herault	The University of Tennessee at Knoxville	2.3.1.09 STPM11 ParSEC: Distributed Tasking	Computer Science	400
CSC250STPM16	Latchesar Alexandrov Lonkov	Los Alamos National Laboratory (LANL)	2.3.1.16 SICM: Simplified Interface to Complex Memory	Computer Science	1,000
CSC250STPM17	Paul Hamilton Hargrove, Erich Strohmaier	Lawrence Berkeley National Laboratory (LBNL)	2.3.1.14 STPM17-UPC++ & GASNet	Computer Science	5,000
CSC250STPR19	Peter Hugh Beckman	Argonne National Laboratory	2.3.1.19 STPR19 Argo: Argo/Power Steering	Computer Science	400
CSC250STPR27	David Edward Bernholdt	Oak Ridge National Laboratory (ORNL)	2.3.1.17 STPR27-OMPI-X: Open MPI for Exascale	Materials Science	4,000
CSC250STTO09	Hartwig Andreas Anzt, Anthony Danalis, Earl Luther Carr, Heike Jagode	The University of Tennessee at Knoxville	2.3.2.06 STTO09 EXAPAPI	Computer Science	4,000
CSC250STTO11	John Michael Mellor-Crummey	Rice University	2.3.2.08 STTO11 HPCToolkit	Computer Science	1,300
CSC251HIHE05	Scott Dov Pakin, Simon David Hammond	Los Alamos National Laboratory (LANL)	2.4.2.01 HIHE05-Analytical Modeling - Hardware Evaluation Working Groups	Computer Science	400
CSC251HISD01	Ryan Charles Prout	Los Alamos National Laboratory (LANL)	2.4.4.01 HISD01-Software Integration	Computer Science	400
CSCSTDT12345	Patrick McCormick	Los Alamos National Laboratory (LANL)	2.3.2.12 Flang: Open-source Fortran Front End for the LLVM Infrastructure	Computer Science	400
CVD_CityCOVID	Jonathan Ozik	Argonne National Laboratory	Agent-based Model Called CityCOVID Capable of Tracking Detailed COVID-19 Transmission	Biological Sciences	49,662

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
darkskyml_aesp	Salman Habib	Argonne National Laboratory	Dark Sky Mining	Physics	8,000
datascience	Venkatram Vishwanath	Argonne National Laboratory	ALCF Data Science and Workflows Allocation	Internal	250,000
dcp35	George Em Karniadakis	Brown University	Quantification of Extreme Weather Events and Their Future Changes Using Physics-Informed DeepONet Modeling and Functional Priors	Mathematics	5,000
DDICF-Dev	Duc Minh Cao	University of Rochester	Direct-Drive Inertial Confinement Fusion Code Porting and Proposal Preparation	Fusion Energy	25,714
DesMultiCat	Rafael Gomez-Bombarelli	Massachusetts Institute of Technology (MIT)	Inverse Design of Multicomponent Oxide Catalysts with Generative Models and DFT	Materials Science	23,385
dist_relational_alg	Sidharth Kumar	The University of Alabama at Birmingham	Distributed Relational Algebra at Scale	Computer Science	9,892
DL_MODEX	Maruti Kumar Mudunuru	Pacific Northwest National Laboratory (PNNL)	Towards a Robust and Scalable Deep Learning Workflow for Fast, Accurate, and Reliable Calibration of Watershed Models	Earth Science	32,708
DNS3D	Ramesh Balakrishnan	Argonne National Laboratory	Direct Numerical Simulation of Three-Dimensional Turbulence	Engineering	10,585
DNSVCMHD	Keith Daniel Brauss	Francis Marion University	DNS Simulations of Velocity-Current Magnetohydrodynamic Equations	Mathematics	5,734
DNS_SV_Turb_2WC	Josin Tom	Duke University	DNS Study of Particle Settling Velocities in Turbulence in the Presence of Two-way Coupling	Engineering	4,225
DynCap	Zhiling Lan	Illinois Institute of Technology (IIT)	Dynamic Power Capping for Scientific Applications	Computer Science	8,181
dynstall_ss	Sarasija Sudharsan	Iowa State University (ISU)	Time-resolved Simulations of Unsteady, Separated Flows	Engineering	8,041
E3SM	Lukasz Dariusz Lacinski	Argonne National Laboratory	Energy Exascale Earth System Model	Earth Science	10,000
Eagle_Testing	Avanthi Madduri	Argonne National Laboratory	Eagle Testing Purposes - Updated	Internal	0

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
EarthWorks	Richard Dana Loft	National Science Foundation	Preparing EarthWorks for GPU-Based Climate Simulations at Global Storm- Resolving Scales	Earth Science	80
ECP_SDK	Sameer Suresh Shende	University of Oregon	Deploying the ECP SDK Software Stack at ALCF	Computer Science	3,414
EE-ECP	Xingfu Wu, Valerie Taylor	Argonne National Laboratory	Energy Efficient Tradeoff Among Execution Time and Power of ECP Applications	Computer Science	14,717
electrolyte-chibueze	Chibueze Vincent Amanchukwu	The University of Chicago (UChicago)	Molecular Dynamics-driven Discovery of Novel Liquid Electrolytes	Energy Technologies	28,672
EngineDNS	Christos Frouzakis	Eidgenössische Technische Hochschule Zürich (ETH Zurich)	Towards Reactive DNS in Complex Internal Combustion Engine Geometries	Engineering	42,569
EstopSim_DD	Yosuke Kanai	The University of North Carolina-Chapel Hill	Massively Parallel Electronic Stopping Simulations of High Energy Particles in Solvated DNA	Chemistry	92,984
FDTD_Cancer_2a	Allen Taflove	Northwestern University	Computational Physical Genomics: Exploring Potential Novel Cancer Therapies	Biological Sciences	30,000
field_scale_modeling	Kaiyu Guan	University of Illinois at Urbana-Champaign	Field-scale Coupled Energy-water-carbon-nutrient Simulation Over Agricultural Landscape in the U.S.	Earth Science	32,000
fusiondl_aesp	William Tang	Princeton University	Computational Design of Polymer Grafted Nanoparticle Membrane	Fusion Energy	1,000
GNPMem	Tarak K. Patra	Indian Institute of Technology Madras	Computational design of polymer grafted nanoparticle membrane	Materials Science	4,116
GrainBoundaries	Wissam A. Saidi	University of Pittsburgh	Structure and Properties of Grain Boundaries in Materials for Energy Applications	Materials Science	64,757
GWrealtimeBSE	Jin Zhao	University of Science and Technology of China	Ab initio Simulation for Exciton Dynamics in Two-dimensional Materials Using GW + Real-time BSE	Materials Science	16,994
HACC_aesp	Katrin Heitmann	Argonne National Laboratory	Extreme-Scale Cosmological Hydrodynamics	Physics	4,000
HARDXBIOCANCEL	Carles Serrat	Universitat Politècnica de Catalunya-Barcelona Tech	Coherent Hard X-ray Core Nonlinear Selective Cancellation of the Effect of Biological Target Molecules in Pathogens	Chemistry	16,368

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
HEPCloud-FNAL	Burt Holzman	Fermi National Accelerator Laboratory (Fermilab)	High Energy Physics Computing for Fermilab Experiments via HEPCloud	Physics	16,095
HEP_on_HPC	Jim B. Kowalkowski, Marcarc Francis Paterno	Fermi National Accelerator Laboratory (Fermilab)	HEP Analysis Workflows HPC	Physics	59,680
HighReyTurb_PostProc	Robert D. Moser, Myoungkyu Lee	The University of Alabama	Data Analysis of Turbulent Channel Flow at High Reynolds Number	Engineering	30,981
HiPressMulticompflow	Hongyuan Zhang	University of Minnesota-Twin Cities	Physics-Based Modeling of Multicomponent Transcritical Phase Change and Spray Breakup in High-Pressure Liquid-Fueled Combustors	Engineering	5,909
HIV-PR	Ao Ma	University of Illinois at Chicago	Understanding the Mechanism of Ligand-induced Conformational Dynamics of HIV-1 Protease and the Effects of Mutations	Biological Sciences	38,914
HNPballistics	Sinan Keten	Northwestern University	Engineering Nanocellulose Based Hairy Nanoparticle Assemblies for High Ballistic Impact Performance	Engineering	32,000
hpc-spectacle	Kevin Antoney Brown	Argonne National Laboratory	Evaluate and Optimize Data Movement Strategies in AI and Climate Science Workloads	Computer Science	20,753
hpc_me	Thomas Edward Robinson	National Oceanic and Atmospheric Administration (NOAA)	HPC Portable Containers for Model Environments	Earth Science	10,000
HumanVAN	Grant Addison Hartung	Athinoula A. Martinos Center for Biomedical Imaging	Biophysical Modeling of the Functional MRI Signal Through Parametric Variations in Neuronal Activation and Blood Vessel Anatomy Using Realistic Synthetic Microvascular…	Biological Sciences	44,581
Intel	Kalyan Kumaran, Scott Parker, Timothy Joe Williams, Venkatram Vishwanath	Argonne National Laboratory	Intel Employees in Support of Theta	Internal	4,000
IntelVis	Joseph A. Insley	Argonne National Laboratory	Intel Visualization Development	Computer Science	26,327

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
JCESR	Larry Curtiss, Anubhav Jain	Argonne National Laboratory	Development of High Throughput Methods	Materials Science	117,184
Job_Interference	Zhiling Lan	Illinois Institute of Technology (IIT)	Workload Interference Analysis on Theta	Computer Science	23,209
LatticeQCD_aesp	Paul Mackenzie, Norman Howard Christ	Fermi National Accelerator Laboratory (Fermilab)	Lattice Quantum Chromodynamics Calculations for Particle and Nuclear Physics	Physics	97,500
LESDNSHTECESHE2021	Lane Benjamin Carasik	Virginia Commonwealth University (VCU)	LES and DNS of Heat Transfer Enhancements in Clean Energy System Heat Exchangers	Engineering	5,083
lhtes	Kedar Prashant Shete	University of Massachusetts-Amherst	DNS of Turbulent Phase Changing Flows	Energy Technologies	5,609
LIGHTCONTROL	Sandra Gail Biedron	University of New Mexico	Light Sources and Their Control Using AI Techniques	Physics	9,247
lipid-sampling	Yun Lyna Luo	Western University of Health Sciences	Development of Enhanced Sampling Approach for Heterogenous Membrane	Biological Sciences	56,186
LoopSynch	Shina Caroline Lynn Kamerlin	Georgia Institute of Technology (Georgia Tech)	Link Between Loop Dynamics and Turnover Number Across 30 Extant Triosephosphate Isomerases	Chemistry	13,899
LQCDdev	James Clifton Osborn	Argonne National Laboratory	Lattice QCD Development	Physics	1,849
lqcdml_aesp	William Detmold	Massachusetts Institute of Technology (MIT)	Machine Learning for Lattice Quantum Chromodynamics	Physics	13,500
LSSMEQ	Marco Govoni	Argonne National Laboratory	Large-scale Simulations of Materials for Energy and Quantum Information Science	Materials Science	5,788
LTC_Aramco_theta	Roberto Torelli	Argonne National Laboratory	Investigation of Gasoline-Range Fuels for a Heavy-Duty Diesel Engine in a Low-Temperature Combustion Regime	Engineering	20,500
magnetotail	Samuel Richard Totorica	Princeton University	Kinetic Simulations of the Dynamic Magnetotail	Physics	18,035
Maintenance	William Edward Allcock, John Francis O'Connell, John Patrick Reddy, Ryan Milner, Torrance Ivan Leggett	Argonne National Laboratory	LCF Operations System Maintenance	Internal	6,245
marine-twin	Flavio Dal Forno Chuahy	Oak Ridge National Laboratory (ORNL)	Marine Digital-Twin Full-Scale Simulation	Energy Technologies	14,619

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
matml_aesp	Noa Marom	Carnegie Mellon University	Many-Body Perturbation Theory Meets Machine Learning to Discover Singlet Fission Materials	Materials Science	165,049
MDClimSim	Veerabhadra Rao Kotamarthi	Argonne National Laboratory	High Resolution Regional Climate Model Simulations	Earth Science	64,000
MEDDIAC	Ehud Strobach	Agricultural Research Organization	High Resolution Interactions of Mediterranean Cyclone with Ocean Eddies	Earth Science	32,000
metastable	Subramanian Sankaranarayanan	Argonne National Laboratory	Metastable Phase Diagram of Material	Materials Science	21,274
MI2Dmaterials	Trevor David Rhone	Rensselaer Polytechnic Institute (RPI)	Materials Informatics Study of Two-dimensional Magnetic Materials and Their Heterostructures	Materials Science	91,700
MKM_catal	Wilfred T. Tjalke Tysoe	University of Wisconsin-Milwaukee	Enantioselectivity in Heterogeneous Catalysts via the Addition of Chiral Modifiers	Chemistry	14,331
ML-Coupling	Shinhoo Kang	Argonne National Laboratory	Data-driven Coupling Methods for Atmospheric-Ocean Interactions	Earth Science	5,920
ML-target	Lianshan Lin	Oak Ridge National Laboratory (ORNL)	Application of ML to the Liquid Mercury Target	Physics	16,071
MLP4THERMO	Cem Sevik	Eskisehir Technical University	Machine Learning Potentials for Thermal Properties of Two-Dimensional Materials	Materials Science	29,412
MOAB_App	Vijay Subramaniam Mahadevan	Argonne National Laboratory	MOAB Algorithmic Performance Portability	Mathematics	37,951
ModelingCoronaVirus	Zhangli Peng	University of Illinois at Chicago	Modeling Corona Virus	Biological Sciences	197,958
MoltenSalts	Nicholas Everett Jackson	University of Illinois at Urbana-Champaign	Automated Active Learning on ALCF for Machine Learning Forcefield Automation	Nuclear Energy	30,000
MPICH_MCS	Kenneth James Raffenetti, Pavan Balaji	Argonne National Laboratory	MPICH - A High-Performance and Widely Portable MPI Implementation	Computer Science	30,154
MPI_Aurora_Intel	Devi Sudheer Kumar Chunduri	Argonne National Laboratory	MPI Development for Aurora	Computer Science	6,000
multimode_comb	Pinaki Pal	Argonne National Laboratory	High-Fidelity CFD Simulations of Multi-Mode Combustion	Energy Technologies	21,831

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
multiphysics_aesp	Amanda Randles	Duke University	Extreme-scale In Situ Visualization and Analysis of Fluid-Structure-Interaction Simulations	Engineering	10,000
MultScale_Frac_DON	Somdatta Goswami	Brown University	A Multiscale Surrogate Model for Fracture Evolution using DeepONet	Engineering	1,000
NAMD_aesp	Benoit Roux, James Christopher Phillips	The University of Chicago (UChicago)	Free Energy Landscapes of Membrane Transport Proteins	Biological Sciences	19,500
NAQMC_RMD_aesp	Aiichiro Nakano	University of Southern California (USC)	Metascalable Layered Materials Genome	Materials Science	10,000
nek52rs	Aleksandr V. Obabko	Argonne National Laboratory	Nek5000/NekRS for NRC and COVID LES	Engineering	8,470
nek5nrc	Aleksandr V. Obabko	Argonne National Laboratory	Nek5000 NRC Support	Engineering	16,000
NekM1Comb	Sicong Wu	Argonne National Laboratory	Development of High-Fidelity Simulations of Gas Turbine Combustors for Sustainable Aviation Applications	Energy Technologies	230,044
Nek_Boost	Pinaki Pal	Argonne National Laboratory	Development of High-Fidelity and Efficient Modeling Capabilities for Enabling Co-Optimization of Fuels and Multi-Mode Engines	Energy Technologies	21,643
networkbench	Kevin Harms, Elise Jennings, Misbah Mubarak	Argonne National Laboratory	Network Benchmarking and Modeling	Computer Science	3,312
neutrino_osc_ADSP	Marco Del Tutto	Fermi National Accelerator Laboratory (Fermilab)	Machine Learning for Data Reconstruction to Accelerate Physics Discoveries in Accelerator-Based Neutrino Oscillation Experiments	Physics	37,206
NextGenReac	Yiqi Yu	Argonne National Laboratory	Toward the Future: High-Fidelity Simulation for Next Generation Nuclear Reactors	Nuclear Energy	47,157
NitrateRemoval	Joshua Jodhimani Gabriel	Argonne National Laboratory	Selective Electrochemical Reduction of Nitrate to Value-Added Products Using a Reactive Electrochemical Membrane System	Chemistry	14,958
novacosmics	Alexander I. Himmel	Fermi National Accelerator Laboratory (Fermilab)	NOvA Cosmic Rejection	Physics	8,683

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
NUCat_Micro-CT	Marta Garcia Martinez	Argonne National Laboratory	Large Volume Feline Spinal Cord Microtomography	Biological Sciences	26,170
NucContainmentMix	Christopher Fred Boyd	U.S. Nuclear Regulatory Commission (NRC)	LES Simulations of Severe Accident Conditions in Nuclear Containment	Nuclear Energy	14,815
NWChemEx_aesp	Theresa Windus, Alvaro Vazquez Mayagoitia	Pacific Northwest National Laboratory (PNNL)	NWChemEx: Tackling Chemical, Materials & Biochemical Challenges in the Exascale Era	Chemistry	2,000
Operations	William Edward Allcock	Argonne National Laboratory	Systems Administration Tasks	Internal	78,125
OptADDN	Sandeep Madireddy	Argonne National Laboratory	Optimal Architecture Discovery for Deep Probabilistic Models and Neuromorphic Systems	Computer Science	30,000
PARTURB3D	Ramesh Balakrishnan	Argonne National Laboratory	Simulating Turbulent Particulate Flows Inside Enclosures	Engineering	190,922
Performance	Scott Parker, Raymond M. Loy	Argonne National Laboratory	Performance	Internal	500,000
PHASTA_aesp	Kenneth Edward Jansen	University of Colorado-Boulder	Extreme Scale Unstructured Adaptive CFD: From Multiphase Flow to Aerodynamic Flow Control	Engineering	19,500
PHASTA_NCSU	Igor A. Bolotnov	North Carolina State University (NCSU)	Multiphase Simulations of Nuclear Reactor Thermal Hydraulics	Engineering	79,170
PodPre	Romit Maulik	Argonne National Laboratory	Predictability Analysis of the ERA5 Dataset	Earth Science	1,187
psr001	Ronald Otis Grover	General Motors Company	Electric Motor Thermal Management Analysis	Engineering	20,946
PTLearnPhoto	Noa Marom	Carnegie Mellon University	Many-Body Perturbation Theory Meets Machine Learning to Discover Materials for Organic Photovoltaics	Materials Science	630,436
PUR-IRL	Nicholas Lee-Ping Chia	Mayo Clinic-Minnesota	Inferring the Reward Function of Cancer	Biological Sciences	8,312
QCProxyApps	Graham Donald Fletcher	Argonne National Laboratory	Quantum Chemistry Proxy Applications	Chemistry	10,388
QMCPACK_aesp	Anouar Benali	Argonne National Laboratory	Extending Moore's Law Computing with Quantum Monte Carlo	Materials Science	19,500

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
qsars_qm_vae	Brad Reisfeld	Colorado State University	Discovering Quantitative Structure Activity Relationships Using Quantum Chemical Descriptors and Variational Autoencoding	Biological Sciences	257
QuantumDS	Alvaro Vazquez Mayagoitia	Argonne National Laboratory	Quantum Mechanics and Data Science	Chemistry	6,458
radix-io	Philip Hutchinson Carns	Argonne National Laboratory	System software to enable data-intensive science	Computer Science	42,696
RAPINS	Eliu Antonio Huerta Escudero	Argonne National Laboratory	Reproducible and Accelerated Physics-inspired Neural Networks	Physics	16,858
RaptorX	Jinbo Xu	Toyota Technological Institute at Chicago (TTIC)	Protein Folding through Deep Learning and Energy Minimization	Biological Sciences	45,554
RCM_4km	Jiali Wang	Argonne National Laboratory	Generation of a Next Level Dataset for Regional Scale Climate Modeling: Convective Resolving Spatial Scales	Earth Science	125,468
Redox_ADSP	Logan Timothy Ward	Argonne National Laboratory	Autonomous Molecular Design for Redox Flow Batteries	Materials Science	13,177
REI_Flares	Marc Cremer	Reaction Engineering International	Leveraging the UCF for Simulation of Industrial Flares	Chemistry	39,363
RL-fold	Arvind Ramanathan	Argonne National Laboratory	Targeting Intrinsically Disordered Proteins Using Artificial Intelligence Driven Molecular Simulations	Biological Sciences	23,182
rnn-robustness	Liam Benjamin Johnston	University of Wisconsin-Madison	Large-scale Factorial Experiment on RNN Robustness	Computer Science	157,232
RTI_DD	Tapan Kumar Sengupta	Indian Institute of Technology Dhanbad	Peta- and Exa-Scale Computing of Rayleigh-Taylor Instability	Engineering	13,747
r_workshop	Paige Carolyn Kinsley, Yasaman Ghadar	Argonne National Laboratory	R Workshop	Training	32
sbi-fair	Pete Beckman, Kamil Antoni Iskra	Argonne National Laboratory	FAIR Surrogate Benchmarks Supporting AI and Simulation Research	Computer Science	6,197
scalablemoose	Fande Kong	Idaho National Laboratory (INL)	MOOSE Scaling Study	Nuclear Energy	2,630
SCPlasma	Ranganathan Gopalakrishnan	University of Memphis	Thermodynamics and Transport Models of Strongly Coupled Dusty Plasmas	Physics	47,005

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
SDL2021x	Wai Nim Alfred Tang	Phd Tutor Hub	SDL2021 Extension	Computer Science	82
SEEr-planning	Zhiling Lan	Illinois Institute of Technology (IIT)	Performance and Power Tradeoff Analysis of AI-Enabled Science on CPU-GPU System	Computer Science	15,157
SeismicHazard_2	Christine Anne Goulet	University of Southern California (USC)	Extreme-Scale Simulations for Advanced Seismic Ground Motion and Hazard Modeling	Earth Science	32,000
SENSEI	Silvio Humberto Rafael Rizzi, Joseph A. Insley, Nicola Joy Ferrier, Venkatram Vishwanath	Argonne National Laboratory	Scalable Analysis Methods and In Situ Infrastructure for Extreme Scale Knowledge Discovery	Computer Science	51,787
Shavalier_sims	Joshua Daniel Gezelter	University of Notre Dame	Heat Transport in Gold Interfaces Capped with Thiolated Oligoethylene Glycol Using a Polarizable Force Field” and “Heat Transport in Gold Interfaces Capped with…	Chemistry	7,034
SolarWindowsADSP	Jacqueline Manina Cole	University of Cambridge	Data-Driven Molecular Engineering of Solar-powered Windows	Materials Science	9,862
SPALL	Mauricio Rene Ponga	The University of British Columbia	Two-pulses Laser Spall Numerical Experiments	Engineering	32,000
spentFuel	Angela Di Fulvio	University of Illinois at Urbana-Champaign	Cask Mis-loads Evaluation Techniques	Nuclear Energy	46,184
startup	John Aaron Palmore	Virginia Polytechnic Institute and State University (Virginia Tech)	High-Fidelity Simulations of Spray and Droplet Combustion (Startup)	Chemistry	8,000
SuperBERT	Ian Foster	Argonne National Laboratory	Training of Language Models on Large Quantities of Scientific Text	Computer Science	29,232
swift	Venkatram Vishwanath	Argonne National Laboratory	Application of High-Performance Computing (HPC) and Artificial Intelligence (AI) to Explore Patterns and Anomalies in Swift Data Streams	Computer Science	2,000
TDMD_thermostat	Charles Michael McCallum	University of the Pacific	Effect of the Thermostat on the Simulation of ESI Processes	Chemistry	1,500
THGSupport	Kevin Harms	Argonne National Laboratory	The HDF Group Support	Computer Science	754

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
TNContract	James Clifton Osborn	Argonne National Laboratory	Tensor Network Contractions for QIS	Physics	11,403
TomoEncoders	Rajkumar Kettimuthu	Argonne National Laboratory	TomoEncoders: Computer Vision Framework for 4D X-ray Tomography	Engineering	2,112
Tools	Scott Parker	Argonne National Laboratory	ALCF Performance Tools	Internal	5,000
TotalView	Peter Michael Thompson, Raymond M. Loy	Rogue Wave Software, Inc.	TotalView Debugger on Blue Gene P	Internal	1,000
TRANScrispr	Giulia Palermo	University of California-Riverside	Dynamics and Mechanism of Transposon-Encoded CRISPR-Cas Systems	Biological Sciences	50,000
TRB	Parisa Mirbod	University of Illinois at Chicago	Turbulent Rayleigh-Benard Convection in Suspensions of Bubbles	Engineering	39,034
UINTAH_aesp	Martin Berzins, John Andrew Schmidt	The University of Utah	Design and Evaluation of High-efficiency Boilers for Energy Production Using a Hierarchical V/UQ Approach	Chemistry	15,500
Ultrafast_X-ray	Jin Wang	Argonne National Laboratory	Ultrafast_X-ray	Engineering	63,435
User_Services	Haritha Siddabathuni Som	Argonne National Laboratory	User Services	Internal	0
VeloC	Bogdan Florin Nicolae	Argonne National Laboratory	VeloC: Very Low Overhead Checkpointing System	Computer Science	63,442
Vendor_Support	William Edward Allcock, Andrew J. Cherry, Susan Marie Coghlan, Torrance Ivan Leggett, William R. Scullin	Argonne National Laboratory	Vendor Support	Internal	31
vib_free_eng	Buu Q. Pham	Ames Laboratory	Vibrational Free Energy from Quantum-chemical Calculations of Large Molecular Systems	Chemistry	15,333
VIPRA	Ashok Srinivasan	University of West Florida	Simulation of Viral Infection Propagation Through Air-Travel	Computer Science	4,885
visualization	Joseph A. Insley, Michael E. Papka	Argonne National Laboratory	Visualization and Analysis Research and Development for ALCF	Internal	25,000
Viz_Support	Joseph A Insley, William Edward Allcock	Argonne National Laboratory	Visualization Support	Computer Science	4,000

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
wall_turb_dd	Ramesh Balakrishnan	Argonne National Laboratory	Wall Resolved Simulations of Canonical Wall Bounded Flows	Engineering	30,928
WaterHammer	Hong Zhang, Hong Zhang	Argonne National Laboratory	Water Hammer Simulation	Mathematics	6,666
XGC_aesp	Choongseok Chang	Princeton Plasma Physics Laboratory (PPPL)	High-Fidelity Simulation of Fusion Reactor Boundary Plasmas	Fusion Energy	19,500
XMultiImage	Phay J. Ho	Argonne National Laboratory	Multimodal Imaging with Intense X-ray Pulses	Chemistry	17,717
				Total DD	6,722,135

January 1, 2022–December 31, 2022
Director's Discretionary (DD) Projects on Thetagpu

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
2D-magnet	Can Ataca	University of Maryland, Baltimore County	Correlated Two-Dimensional Magnets at Chemical Accuracy	Materials Science	2,000
ACO2RDS	John J. Low	Argonne National Laboratory	Adsorptive CO2 Removal from Dilute Sources	Materials Science	18,823
AF2C-high-mem	Ada Anna Sedova	Oak Ridge National Laboratory (ORNL)	Large Model Inference for Deep Neural Network Models of Protein Complexes	Biological Sciences	2,000
AGI-for-Science	Rick Lyndon Stevens	Argonne National Laboratory	Large scale multi modal language models for science comprehension	Computer Science	398
AI-based-NDI-Spirit	Rajkumar Kettimuthu	Argonne National Laboratory	Framework and Tool for Artificial Intelligence & Machine Learning Enabled Automated Non-Destructive Inspection of Composite Aerostructures Manufacturing	Engineering	1,941
AI-Steer	Rajkumar Kettimuthu	Argonne National Laboratory	AI-Steer: AI-driven online steering of light source experiments	Computer Science	2,000
AI4NMR	Eric Michael Jonas	The University of Chicago (UChicago)	Structure Elucidation for Nuclear Magnetic Resonance via Structured Prediction	Chemistry	301
AIASMAAR	Rui Hu	Argonne National Laboratory	Artificial Intelligence Assisted Safety Modeling and Analysis of Advanced Reactors	Nuclear Energy	1,610
AIHacks	Eliu Antonio Huerta Escudero	Argonne National Laboratory	AI and HPC Applications for Experimental Science	Materials Science	2,000
AI_Acceleration	Ian Foster	Argonne National Laboratory	Exploration of AI Accelerators for Neural Networks	Computer Science	258
ALCFAITP	Venkatram Vishwanath	Argonne National Laboratory	Argonne AI Training Program	Training	7,009
alphafold	Michael E. Papka	Argonne National Laboratory	AlphaFold as a Community Service	Biological Sciences	10,000
AP4GPU	Kevin Harms	Argonne National Laboratory	AutoPerf for Nvidia GPU	Computer Science	250
APSDataAnalysis	Rafael Vescovi	Argonne National Laboratory	APS Beamline Data Processing and Analysis	Computer Science	2,521
ArgonneGPUAccess	Craig Stacey	Argonne National Laboratory	GPU Access for Argonne Researchers	Internal	2,142

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
ATPESC2022	Raymond M. Loy	Argonne National Laboratory	Argonne Training Program for Extreme-Scale Computing 2022	Computer Science	2,000
ATPESC_Instructors	Raymond M. Loy	Argonne National Laboratory	Argonne Training Program on Extreme-Scale Computing for ALL Instructors	Training	500
AutoPhase	Yudong Yao	Argonne National Laboratory	Real-time X-ray Coherent Imaging with a Self-trained Neural Network	Computer Science	1,514
autopology_alcf	Rafael Gomez-Bombarelli	Massachusetts Institute of Technology (MIT)	End to End Classical Force Field Parametrization for Polymer Electrolytes Using Machine Learning	Materials Science	2,723
BES-AXMAS	Charlotte Lisa Haley	Argonne National Laboratory	AI for Quality Control of Single Crystal X-ray Scattering Experiments	Materials Science	526
biolearning	Chongle Pan	University of Oklahoma	Development of Large-scale Biomedical Machine Learning Models	Biological Sciences	1,329
BioMed	Yuri Alexeev	Argonne National Laboratory	Research Hypotheses from Existing Biomedical Papers	Biological Sciences	1,846
BIP167	Philip Kurian	Howard University	Computing Superradiance and van der Waals Many-body Dispersion Effects for Biomacromolecules	Physics	1,341
BirdAudio	Nicola Joy Ferrier	Argonne National Laboratory	Machine Learning for Classificaiton of Birdsong	Computer Science	2,299
BLawB	Lucian Ivan	Canadian Nuclear Laboratories	Application of Maximum-Entropy Moment Methods to Turbulent and Multiphase Flow Prediction: Software Package Preparation	Engineering	300
BlazingSQLforHPC	Benjamin Hernandez Arreguin	Oak Ridge National Laboratory (ORNL)	Optimizing BlazingSQL for DOE's Leadership Computing Facilities	Computer Science	258
bloodflow_dd	Jifu Tan	Northern Illinois University (NIU)	Multiphysics Modeling of Biological Flow with Cell Suspensions	Engineering	2,483
BNN-Scale	Murali Krishna Emani	Argonne National Laboratory	Optimizing Bayesian Neural Networks for Scientific Machine Learning Applications	Computer Science	1,759
BPC	Christopher Michael Graziul	The University of Chicago (UChicago)	Optimization of Audio Processing Pipeline for Broadcast Police Communications	Computer Science	1,117
BRAIN	Getnet Dubale Betrie	Argonne National Laboratory	Scalable Brain Simulator for Extreme Computing	Biological Sciences	2,986

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
BS-SOLCTRA	Esteban Meneses	Costa Rica National High Technology Center	Plasma Physics Simulations for SCR-1 Stellarator	Physics	2,472
bubble-ai	Ben J. Blaiszik	Argonne National Laboratory	Discovery of Novel Fuel Cell Catalyst Materials via Development of High-Throughput AI-Guided Characterization Methods	Materials Science	230
CAIDS	Julio Cesar Mendez Carvajal	North Carolina State University (NCSU)	Consistent Averaging Procedure for Solving the Fundamental Equations of Fluid Dynamics	Engineering	422
candle_aesp	Rick Lyndon Stevens, Thomas Scott Brettin	Argonne National Laboratory	Virtual Drug Response Prediction	Biological Sciences	250
Carbon_composites	Hendrik Heinz	University of Colorado-Boulder	Designing Functional Nanostructures and Carbon-Based Composite Materials	Materials Science	590
Catalyst	Katherine M. Riley, Christopher James Knight, James Clifton Osborn, Timothy Joe Williams	Argonne National Laboratory	Catalyst	Internal	4,000
ceed-app	Misun Min	Argonne National Laboratory	Aerosoal Transport Modeling Towards Exascale	Engineering	3,989
CeleritasA100	Paul Kollath Romano	Argonne National Laboratory	Performance Measurements of Celeritas Monte Carlo transport for HEP	Physics	200
cfdml_aesp	Kenneth Edward Jansen	University of Colorado-Boulder	Data Analytics and Machine Learning for Exascale CFD	Engineering	25
CharmRTS	Laxmikant Kale, Abhinav Bhatele, Juan Jose Galvez-Garcia	University of Illinois at Urbana-Champaign	Charm++ and Its Applications	Computer Science	2,000
Clouds	Ian Foster	Argonne National Laboratory	Unsupervised Analysis of Satellite Cloud Imagery	Earth Science	8,628
compsensingADSP	Robert Hovden	University of Michigan	Dynamic Compressed Sensing for Real-time Tomographic Reconstruction	Materials Science	1,315
Comp_Perf_Workshop	Raymond M. Loy, Yasaman Ghadar	Argonne National Laboratory	ALCF Computational Performance Workshop	Training	1,000
connectomics_aesp	Nicola Joy Ferrier, Thomas David Uram	Argonne National Laboratory	Enabling Connectomics at Exascale to Facilitate Discoveries in Neuroscience	Biological Sciences	822

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
coreform-lattice	Greg Vernon	Coreform LLC	HPC-enabled Geometry-compliant Lattice Structures for 3D Printing and Structural Simulation	Engineering	1,527
covid-ct	Ravi Kiran Madduri	Argonne National Laboratory	Medical Imaging Domain-Expertise Machine Learning for Interrogation of COVID	Computer Science	3,635
cray-hpo	Michael Adnan Salim	Argonne National Laboratory	Scaling Studies of CrayAI Hyperparameter Optimization	Computer Science	200
crocus	Dimitrios Fytanidis	Argonne National Laboratory	NekRS Scaling Urban Modeling -DOE BER CROCUS Proposal	Earth Science	1,000
CSC249ADOA01	Rick Lyndon Stevens, Thomas Scott Brettin	Argonne National Laboratory	2.2.4.03 ADOA01 CANDLE: Exascale Deep Learning Enabled Precision Medicine for Cancer	Biological Sciences	4,037
CSC250STDM14	Franck Cappello	Argonne National Laboratory	2.3.4.14 STDM14 - VeloC-SZ: Very Low Overhead Transparent Multilevel Checkpoint/Restart/SZ: Fast, Effective, Parallel Error-bounded Exascale Loss....	Computer Science	1,000
CVD_CityCOVID	Jonathan Ozik	Argonne National Laboratory	Agent-based Model Called CityCOVID Capable of Tracking Detailed COVID-19 Transmission	Biological Sciences	1,988
datascience	Venkatram Vishwanath	Argonne National Laboratory	ALCF Data Science and Workflows Allocation	Internal	12,500
dcp35	George Em Karniadakis	Brown University	Quantification of Extreme Weather Events and Their Future Changes Using Physics-Informed DeepONet Modeling and Functional Priors	Mathematics	2,000
Deep_WF	Zhi Qiao	Argonne National Laboratory	AI-enabled Real-time Super-resolution X-ray Wavefront Sensing and Advanced Beamline Control	Physics	1,936
dendritesegmentation	Ben J. Blaiszik	Argonne National Laboratory	Machine Learning for Automated Dendrite Segmentation to Accelerate Experiments at the Advanced Photon Source	Materials Science	61
dist_relational_alg	Sidharth Kumar	The University of Alabama at Birmingham	Distributed Relational Algebra at Scale	Computer Science	764

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
DL4VIS	Hanqi Guo	Argonne National Laboratory	Deep Learning for In Situ Analysis and Visualization	Computer Science	3,170
DLHMC	Sam Alfred Foreman	Argonne National Laboratory	Deep Learning HMC	Physics	2,559
DL_MODEX	Maruti Kumar Mudunuru	Pacific Northwest National Laboratory (PNNL)	Towards a Robust and Scalable Deep Learning Workflow for Fast, Accurate, and Reliable Calibration of Watershed Models	Earth Science	1,325
DNS3D	Ramesh Balakrishnan	Argonne National Laboratory	Direct Numerical Simulation of Three-Dimensional Turbulence	Engineering	2,000
DNS_SV_Turb_2WC	Josin Tom	Duke University	DNS Study of Particle Settling Velocities in Turbulence in the Presence of Two-way Coupling	Engineering	422
DPCPPA100	Kevin Harms	Argonne National Laboratory	DPC++ on A100	Computer Science	522
Drug_FEP_Data	Wei Jiang	Argonne National Laboratory	Machine Learning of Drug Binding and Toxicity Based on High Throughput Free Energy Computations	Biological Sciences	1,658
DynamicCS	Jonathan Tyler Schwartz, Huihuo Zheng	Argonne National Laboratory	Dynamic Compressed Sensing for Real-time Tomographic Reconstruction	Materials Science	2,500
DynCatalysis	Anastassia N. Alexandrova	University of California-Los Angeles	Heterogeneous Catalysis as a Collective Phenomenon Within Dynamic Ensembles of States	Chemistry	1,200
E3SM	Lukasz Dariusz Lacinski	Argonne National Laboratory	Energy Exascale Earth System Model	Earth Science	2,000
EarthWorks	Richard Dana Loft	National Science Foundation	Preparing EarthWorks for GPU-Based Climate Simulations at Global Storm- Resolving Scales	Earth Science	256
ECP_SDK	Sameer Suresh Shende	University of Oregon	Deploying the ECP SDK Software Stack at ALCF	Computer Science	3,942
EE-ECP	Xingfu Wu, Valerie Taylor	Argonne National Laboratory	Energy Efficient Tradeoff Among Execution Time and Power of ECP Applications	Computer Science	1,848
Emerging_Tech_ML	Benjamin Alan Blakely	Argonne National Laboratory	Graph-based Bibliometric Analysis for Emerging Technology Discovery	Computer Science	3,021

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
EngineDNS	Christos Frouzakis	Eidgenössische Technische Hochschule Zürich (ETH Zurich)	Towards Reactive DNS in Complex Internal Combustion Engine Geometries	Engineering	1,951
ES_AI	Himanshu Sharma	Pacific Northwest National Laboratory (PNNL)	Deep Neural Network Model for Modeling Aqueous Aerosol Chemistry for Climate Science	Earth Science	541
EvalDL	Natalia Sergeyevna Vasileva	Cerebras Systems Inc.	Evaluation of Deep Learning Models	Computer Science	8
EVITA	George K. Thiruvathukal	Loyola University Chicago	Energy-efficient VIsual Transformer Architecture: Transformer Models for Deployment on Embedded Systems	Computer Science	988
ExaPF	Michel Schanen	Argonne National Laboratory	Optimal Power Flow on GPUs	Mathematics	48
ExaSGD	Michel Schanen	Argonne National Laboratory	CP ExaSGD ADSE22	Energy Technologies	100
fibregpu	Davide Di Giusto	University of Udine	DNS of Fibre-laden Turbulent Channel Flow	Engineering	595
field_scale_modeling	Kaiyu Guan	University of Illinois at Urbana-Champaign	Field-scale Coupled Energy-Water-Carbon-Nutrient Simulation Over Agricultural Landscape in the U.S.	Earth Science	2,000
Fornax_GPU	Adam Seth Burrows	Princeton University	Porting the Fornax Supernova Code to GPUs	Physics	1,463
fuelspray	Miaoqi Chu	Argonne National Laboratory	Ultrafast X-ray Vision of Fuel Injection and Near-field Spray	Energy Technologies	500
gccy3	Marcarc Francis Paterno	Fermi National Accelerator Laboratory (Fermilab)	Cosmological Parameter Inference from Galaxy Clusters	Physics	1,936
GNN-internship	Eliu Antonio Huerta Escudero	Argonne National Laboratory	Graph Neural Net (Internship)	Computer Science	100
gnn_uq	Shengli Jiang	University of Wisconsin-Madison	Molecular Property Uncertainty Quantification via Automated Graph Neural Networks	Computer Science	791
GNPMem	Tarak K. Patra	Indian Institute of Technology Madras	Computational Design of Polymer Grafted Nanoparticle Membrane	Materials Science	646
GPU-DG	Pinaki Pal	Argonne National Laboratory	GPU-enabled Discontinuous Galerkin Simulations of Complex Fluid Flows	Engineering	1,333

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
gpu_hack	Yasaman Ghadar, Raymond M. Loy	Argonne National Laboratory	GPU Hackathon	Training	1,200
GRACE	Sayan Ghosh	Pacific Northwest National Laboratory (PNNL)	Graph Analytics Codesign on GPUs	Computer Science	2,346
GrainBoundaries	Wissam A. Saidi	University of Pittsburgh	Structure and Properties of Grain Boundaries in Materials for Energy Applications	Materials Science	1,525
graphs_	Dossay Oryspayev	Brookhaven National Laboratory (BNL)	Exploration of Parallelization of Graph Algorithms	Computer Science	1,062
HighReyTurb_PostProc	Robert D. Moser, Myoungkyu Lee	The University of Alabama	Data Analysis of Turbulent Channel Flow at High Reynolds Number	Engineering	1,000
hno1	Maarten de Hoop	Rice University	Learning the Wave Speed and Source to Solution Map for the Helmholtz Equation in Dimension 3	Mathematics	2,200
HNPballistics	Sinan Keten	Northwestern University	Engineering Nanocellulose Based Hairy Nanoparticle Assemblies for High Ballistic Impact Performance	Engineering	2,000
hp-ptycho	Tekin Bicer	Argonne National Laboratory	High Performance 3D Ptychographic Reconstruction and Image Enhancement	Materials Science	1,542
hpc-spectacle	Kevin Antoney Brown	Argonne National Laboratory	Evaluate and Optimize Data Movement Strategies in AI and Climate Science Workloads	Computer Science	8,000
hpc_me	Thomas Edward Robinson	National Oceanic and Atmospheric Administration (NOAA)	HPC Portable Containers for Model Environments	Earth Science	500
IBM-GSS	Venkatram Vishwanath	Argonne National Laboratory	IBM GeoSpatial Software System	Earth Science	1,377
IMEXLBM	Saumil Sudhir Patel	Argonne National Laboratory	ECP ProxyApp Development for the Lattice Boltzmann Method	Computer Science	1,264
img_ai_exp	Venkatram Vishwanath	Argonne National Laboratory	Benchmarking Xray and EM Based Deep Learning AI Expedition Applications	Physics	1,852
Intel	Kalyan Kumaran, Scott Parker, Timothy Joe Williams, Venkatram Vishwanath	Argonne National Laboratory	Intel Employees in Support of Theta	Internal	4,000

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
IonTransES-ML	Boris Kozinsky	Harvard University	IonTransES - Machine Learning of Quantum Energies	Energy Technologies	5,760
JCESR	Larry Curtiss, Anubhav Jain	Argonne National Laboratory	Development of High Throughput Methods	Materials Science	6,440
LAR-EM	Zheng Zhang	The University of Chicago (UChicago)	Electron Tomography from Limited-Angular-Range Data	Materials Science	888
les120	George Em Karniadakis	Brown University	Learning the Sub-Grid Model in Large Eddy Simulations Using Domain-Decomposition Based Parallel Physics-Informed Neural Networks (PINNs)	Mathematics	3,220
LIGHTCONTROL	Sandra Gail Biedron	University of New Mexico	Light Sources and Their Control Using AI Techniques	Physics	2,002
lipid-sampling	Yun Lyna Luo	Western University of Health Sciences	Development of Enhanced Sampling Approach for Heterogenous Membrane	Biological Sciences	2,000
LoopSynch	Shina Caroline Lynn Kamerlin	Georgia Institute of Technology (Georgia Tech)	Link Between Loop Dynamics and Turnover Number Across 30 Extant Triosephosphate Isomerases	Chemistry	171
LQCD-ML	Wai Nim Alfred Tang	Phd Tutor Hub	LQCD parametric regression by neural networks	Physics	1,503
LQCDdev	James Clifton Osborn	Argonne National Laboratory	Lattice QCD Development	Physics	1,008
lqcdml_aesp	William Detmold	Massachusetts Institute of Technology (MIT)	Machine Learning for Lattice Quantum Chromodynamics	Physics	250
M2RL	Sami Khairy	Argonne National Laboratory	Model-based Multi-Fidelity Reinforcement Learning for Neural Architecture Search	Computer Science	2,000
M4DA	Zichao Di	Argonne National Laboratory	Enabling Large-scale Multimodal Data Analysis for the APS	Materials Science	957
MAB-ALLOYS	Deniz Cakir	University of North Dakota	Discovery of Novel Transition Metal **B**oride Alloys	Materials Science	2,000
Maintenance	William Edward Allcock, John Francis O'Connell, John Patrick Reddy, Ryan Milner, Torrance Ivan Leggett	Argonne National Laboratory	LCF Operations System Maintenance	Internal	3,232

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
matml_aesp	Noa Marom	Carnegie Mellon University	Many-Body Perturbation Theory Meets Machine Learning to Discover Singlet Fission Materials	Materials Science	329
metastable	Subramanian Sankaranarayanan	Argonne National Laboratory	Metastable Phase Diagram of Material	Materials Science	1,090
MI2Dmaterials	Trevor David Rhone	Rensselaer Polytechnic Institute (RPI)	Materials Informatics Study of Two-Dimensional Magnetic Materials and Their Heterostructures	Materials Science	1,081
MILC_GPU	Steven Arthur Gottlieb	Indiana University (IU)	Exploring the Muon Anomalous Magnetic Moment Using ThetaGPU Nodes	Physics	1,407
ML-Coupling	Shinhoo Kang	Argonne National Laboratory	Data-driven Coupling Methods for Atmospheric-Ocean Interactions	Earth Science	1,214
MLP4THERMO	Cem Sevik	Eskisehir Technical University	Machine Learning Potentials for Thermal Properties of Two-Dimensional Materials	Materials Science	414
MLPerf_Storage	Huihuo Zheng	Argonne National Laboratory	MLPerf Storage and I/O Benchmarks for Deep Learning	Computer Science	2,000
MOAB_App	Vijay Subramaniam Mahadevan	Argonne National Laboratory	MOAB Algorithmic Performance Portability	Mathematics	1,595
MPICH_MCS	Kenneth James Raffenetti, Pavan Balaji	Argonne National Laboratory	MPICH - A High Performance and Widely Portable MPI Implementation	Computer Science	1,000
mpi_partitioned	Ahmad Afsahi	Queen's University	GPU-Initiated MPI Partitioned Point-to-Point Communication Over NVSHMEM	Computer Science	267
MultiActiveAI	Dario Dematties	Northwestern Argonne Institute of Science and Engineering (NAISE)	Multimodal Intelligence for Federated Edge Computing Simulations	Computer Science	1,504
multiphysics_aesp	Amanda Randles	Duke University	Extreme-scale In Situ Visualization and Analysis of Fluid-Structure-Interaction Simulations	Engineering	447
MultScale_Frac_DON	Somdatta Goswami	Brown University	A Multiscale Surrogate Model for Fracture Evolution using DeepONet	Engineering	1,000
MVAPICH2	Dhabaleswar Kumar Panda	The Ohio State University	Optimizing and Tuning MVAPICH2-GDR Library and Study Its Impact on HPC and AI Applications	Computer Science	3,966

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
nek52rs	Aleksandr V. Obabko	Argonne National Laboratory	Nek5000/NekRS for NRC and COVID LES	Engineering	529
nekMITLL	Aleksandr V. Obabko	Argonne National Laboratory	Porting Nek5000 MITLL Benchmark into NekRS	Engineering	1,000
neutrino_osc_ADSP	Marco Del Tutto	Fermi National Accelerator Laboratory (Fermilab)	Machine Learning for Data Reconstruction to Accelerate Physics Discoveries in Accelerator-Based Neutrino Oscillation Experiments	Physics	1,033
NNM	Yongchao Yang	Michigan Technological University	Deep Learning for Strongly Nonlinear Dynamical Systems	Engineering	971
novacosmics	Alexander I. Himmel	Fermi National Accelerator Laboratory (Fermilab)	NOvA Cosmic Rejection	Physics	2,232
NSCS	Tanwi Mallick	Argonne National Laboratory	Towards Neighborhood Scale Climate Simulations Using AI and Accelerated GPUs	Computer Science	1,885
NUCat_Micro-CT	Marta Garcia Martinez	Argonne National Laboratory	Large Volume Feline Spinal Cord Microtomography	Biological Sciences	2,000
NuQMC	Alessandro Lovato	Argonne National Laboratory	Nuclear Quantum Monte Carlo	Physics	2,000
Operations	William Edward Allcock	Argonne National Laboratory	Systems Administration Tasks	Internal	1,000
OptADDN	Sandeep Madireddy	Argonne National Laboratory	Optimal Architecture Discovery for Deep Probabilistic Models and Neuromorphic Systems	Computer Science	2,000
PARTURB3D	Ramesh Balakrishnan	Argonne National Laboratory	Simulating Turbulent Particulate Flows Inside Enclosures	Engineering	2,234
PDE_ML	Ramin Baghgar Bostanabad	University of California-Irvine	Self-supervised Coupling of Deep Operator Surrogates for Scalable and Transferable Learning	Engineering	3,457
Performance	Scott Parker, Raymond M. Loy	Argonne National Laboratory	Performance	Internal	1,000
proxima	Logan Timothy Ward	Argonne National Laboratory	Proxima-MD: Accelerating Melting Point Computations with Machine Learning	Materials Science	500
PUR-IRL	Nicholas Lee-Ping Chia	Mayo Clinic-Minnesota	Inferring the Reward Function of Cancer	Biological Sciences	1,491
QCProxyApps	Graham Donald Fletcher	Argonne National Laboratory	Quantum Chemistry Proxy Applications	Chemistry	441

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
qsars_qm_vae	Brad Reisfeld	Colorado State University	Discovering Quantitative Structure Activity Relationships Using Quantum Chemical Descriptors and Variational Autoencoding	Biological Sciences	236
QuantumDS	Alvaro Vazquez Mayagoitia	Argonne National Laboratory	Quantum Mechanics and Data Science	Chemistry	1,637
radix-io	Philip Hutchinson Carns	Argonne National Laboratory	System Software to Enable Data-Intensive Science	Computer Science	93
RAPINS	Eliu Antonio Huerta Escudero	Argonne National Laboratory	Reproducible and Accelerated Physics-inspired Neural Networks	Physics	1,982
RaptorX	Jinbo Xu	Toyota Technological Institute at Chicago (TTIC)	Protein Folding through Deep Learning and Energy Minimization	Biological Sciences	4,605
Redox_ADSP	Logan Timothy Ward	Argonne National Laboratory	Autonomous Molecular Design for Redox Flow Batteries	Materials Science	180
remote_offloading	Jose Manuel Monsalve Diaz	Argonne National Laboratory	Exploring Collective Operations with Remote Offloading	Computer Science	1,181
RL-fold	Arvind Ramanathan	Argonne National Laboratory	Targeting Intrinsically Disordered Proteins Using Artificial Intelligence Driven Molecular Simulations	Biological Sciences	29,328
rnn-robustness	Liam Benjamin Johnston	University of Wisconsin-Madison	Large-scale Factorial Experiment on RNN Robustness	Computer Science	3,782
sbi-fair	Pete Beckman, Kamil Antoni Iskra	Argonne National Laboratory	FAIR Surrogate Benchmarks Supporting AI and Simulation Research	Computer Science	3,703
schwinger_bosons	Yaroslav Tserkovnyak	University of California-Los Angeles	Numerical Solution to Schwinger Boson Mean Field Theory	Physics	100
SCPlasma	Ranganathan Gopalakrishnan	University of Memphis	Thermodynamics and Transport Models of Strongly Coupled Dusty Plasmas	Physics	2,000
SCREAM_Calib	Jiali Wang	Argonne National Laboratory	Towards Neighborhood Scale Climate Simulations using AI and Accelerated GPUs	Earth Science	1,029
SDL2021x	Wai Nim Alfred Tang	Phd Tutor Hub	SDL2021 Extension	Computer Science	82
SDL_Workshop	Yasaman Ghadar	Argonne National Laboratory	ALCF Simulation, Data, and Learning Workshop	Training	500

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
SEEr-planning	Zhiling Lan	Illinois Institute of Technology (IIT)	Performance and Power Tradeoff Analysis of AI-Enabled Science on CPU-GPU System	Computer Science	16,491
SeismicHazard_2	Christine Anne Goulet	University of Southern California (USC)	Extreme-Scale Simulations for Advanced Seismic Ground Motion and Hazard Modeling	Earth Science	2,000
SENSEI	Silvio Humberto Rafael Rizzi, Joseph A. Insley, Nicola Joy Ferrier, Venkatram Vishwanath	Argonne National Laboratory	Scalable Analysis Methods and In Situ Infrastructure for Extreme Scale Knowledge Discovery	Computer Science	1,067
skysurvey_adsp	George Frazer Stein	Lawrence Berkeley National Laboratory (LBNL)	Learning Optimal Image Representations for Current and Future Sky Surveys	Physics	657
smarthpc	Rajkumar Kettimuthu	Argonne National Laboratory	System Level Approach to Optimize Neural Architecture Search in HPC Environments	Computer Science	2,920
SolarWindowsADSP	Jacqueline Manina Cole	University of Cambridge	Data-Driven Molecular Engineering of Solar-powered Windows	Materials Science	6,340
SOLLVE	Sunita Chandrasekaran	Brookhaven National Laboratory (BNL)	Scaling OpenMP with LLVm for Exascale Performance and Portability	Computer Science	3,602
spentFuel	Angela Di Fulvio	University of Illinois at Urbana-Champaign	Cask Mis-loads Evaluation Techniques	Nuclear Energy	1,010
SuperBERT	Ian Foster	Argonne National Laboratory	Training of Language Models on Large Quantities of Scientific Text	Computer Science	9,212
swift	Venkatram Vishwanath	Argonne National Laboratory	Application of High-Performance Computing (HPC) and Artificial Intelligence (AI) to Explore Patterns and Anomalies in Swift Data Streams	Computer Science	2,000
SWIFT_CRADA	Venkatram Vishwanath	Argonne National Laboratory	Anomaly Detection for Swift Transaction Streams	Computer Science	1,209
TDMD_thermostat	Charles Michael McCallum	University of the Pacific	Effect of the Thermostat on the Simulation of ESI Processes	Chemistry	750
THGSupport	Kevin Harms	Argonne National Laboratory	The HDF Group Support	Computer Science	754
TMEM_DEL	Diomedes Elias Logothetis	Northeastern University	Molecular Dynamics Simulation on TMEM16A Chloride Channel	Biological Sciences	3,190

Project Name	PI Name	PI Institution	Project Title	Science Field (Short)	Allocation Amount
TomoEncoders	Rajkumar Kettimuthu	Argonne National Laboratory	TomoEncoders: Computer Vision Framework for 4D X-ray Tomography	Engineering	845
Tools	Scott Parker	Argonne National Laboratory	ALCF Performance Tools	Internal	500
TRANScrispr	Giulia Palermo	University of California-Riverside	Dynamics and Mechanism of Transposon-Encoded CRISPR-Cas Systems	Biological Sciences	2,000
TurboPINNs	Paris Perdikaris	University of Pennsylvania (UPenn)	Parallel Simulations of Turbulent Flows with Physics-Informed Neural Networks	Engineering	2,000
Ultrafast_X-ray	Jin Wang	Argonne National Laboratory	Ultrafast_X-ray	Engineering	4,000
User_Services	Haritha Siddabathuni Som	Argonne National Laboratory	User Services	Internal	55
VASPDEK	David Eugene Keller	University of Rochester	Vasp 2.0.1 GPU Timing Simulations	Physics	0
VeloC	Bogdan Florin Nicolae	Argonne National Laboratory	VeloC: Very Low Overhead Checkpointing System	Computer Science	2,478
visualization	Joseph A. Insley, Michael E. Papka	Argonne National Laboratory	Visualization and Analysis Research and Development for ALCF	Internal	1,000
wall_turb_dd	Ramesh Balakrishnan	Argonne National Laboratory	Wall Resolved Simulations of Canonical Wall Bounded Flows	Engineering	1,789
wereszczynski	Jeffery Michael Wereszczynski	Illinois Institute of Technology (IIT)	MD Simulations of Chromatin Modification and Gene Regulation Mechanism	Biological Sciences	1,121
WRFGPU	Veerabhadra Rao Kotamarthi	Argonne National Laboratory	WRF GPU testing	Earth Science	2,079
XMultiImage	Phay J. Ho	Argonne National Laboratory	Multimodal Imaging with Intense X-ray Pulses	Chemistry	494
				Total DD	394,697

Appendix C – ALCF CY2022 Science Highlights

The following Science Highlights were submitted to ASCR for the 2022 OAR performance period.

Rendezvous Algorithms for Large-Scale Modeling and Simulation

The Science

Scientific simulation methods, when implemented in parallel on distributed-memory computers, use one or more communication patterns to send and receive data between processors. Rendezvous communication is a communication pattern that can be useful when processors neither know which other processors to send their data to, nor which processors will be sending them data. Two particle simulation codes – LAMMPS and SPARTA – encountered performance bottlenecks when certain operations that have this communication pattern were run at large scale using more brute-force algorithms. In this study, a general abstraction of rendezvous algorithms was employed, where the rendezvous decomposition need not represent an alternative spatial decomposition of the data.

The Impact

This work demonstrates that rendezvous methods can help reduce performance bottlenecks for a variety of computational tasks performed in particle- and grid-based codes when simpler algorithms do not scale well. The rendezvous algorithms implemented within LAMMPS and SPARTA performed dramatically faster at scale enabling LAMMPS to enumerate bond topologies for molecular systems with billions of atoms and SPARTA to compute grid/surface intersections in models with billions of grid cells and millions of surface elements much more efficiently. These results show that the methods can scale effectively to DOE’s upcoming exascale machines.

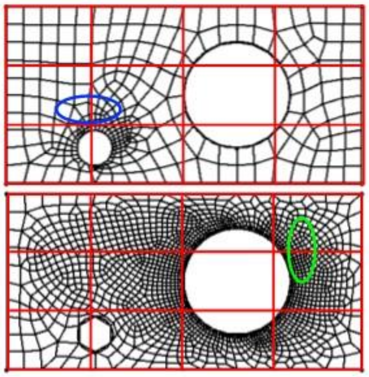

Intermediate rendezvous decomposition for 12 processors (red lines) overlaid on thermal and stress grids. The blue/green colored ovals are clumps of grid cells the same processor owns in the two grids. (Image courtesy of Christopher Knight (ANL) and Steven Plimpton (SNL).

Contact PI: Steven Plimpton (SNL)
ASCR Allocation PI: Christopher Knight (ANL)
ASCR Program/Facility: DD/ALCF
ASCR PM: Saswata Hier-Majumder
Date submitted to ASCR: Jan. 19, 2022
Publication: S.J. Plimpton and C. Knight, *J. Parallel Distr. Com.*, **147**, 184 (2021)

Sandia National Laboratories and Argonne National Laboratory

U.S. DEPARTMENT OF ENERGY | Office of Science

ECP, ALCF

Next-Generation Nonwovens Manufacturing: A Model-Driven Simulation and Machine Learning Approach

The Science

- Developed computational fluid dynamics (CFD) model using Large Eddy Simulations and Volume-Of-Fluids (VOF) to understand air flow behavior and the liquid/gas interface dynamics.
- Explored varying geometry and operating conditions (selected using active machine learning approaches) using the CFD model.
- Identified process-relevant success metrics to optimize using the collected the data.
- Employed a Bayesian optimization strategy (with a Gaussian process surrogate model) to optimize the process and iteratively recommend simulation parameters to test, to reach desired configuration with reduced energy consumption.

The Impact

- Non-woven materials prepared via melt blowing are widely used to make filters (e.g., N95 masks), fabrics, and insulation materials.
- This project aims to reduce the process energy consumption without compromising quality through a coordinated campaign of experiments, simulations, and machine learning.

Goal: Optimize energy efficiency for non-woven applications (e.g., masks, filters, mats), which comprise a 300,000-ton material market, while maintaining process quality. ***Seeking 20% reduction in overall energy consumption.*** This level of reduction in energy consumption would have a global impact for 3M and other manufacturers would likely follow suit.

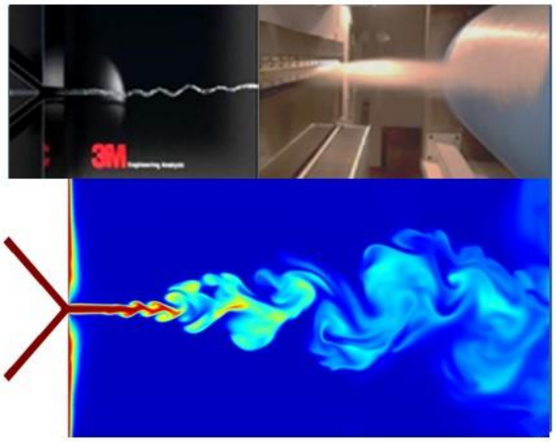

(Top) Melt blowing experiments at 3M and (bottom) Argonne CFD simulation results modeling the same process using conditions selected from machine learning. (Image courtesy of the Argonne and 3M team members on project.)

Contact PI: Ian Foster
ASCR Allocation PI: Ian Foster
ASCR Program/Facility: DD/ALCF
ASCR PM: Saswata Hier-Majumder
Date submitted to ASCR: February 16, 2022

Designing Catalysts under Realistic Operating Conditions from the Atomic Scale

The Science

The oxygen reduction reaction (ORR), a fundamental electrochemical reaction that reduces oxygen molecules to water, is essential to many applications, including hydrogen fuel cells. Platinum is a promising catalyst for ORR in fuel cells, but improvements in catalyst composition and performance are needed to make them a more cost-effective option. However, progress is hampered by trial-and-search experiments due to multiple challenges, including lack of computational tools with required accuracy. To better understand and model the ORR process, researchers developed new computational tools and used highly accurate molecular simulations to predict the catalytic activity of several platinum nanostructures.

The Impact

The team's computational methods can be used to design nanostructures that maximize catalytic efficiency, as well as possible surface modifications to further optimize the cost-benefit ratio of fuel cells. The team's simulations of electrode-electrolyte interfaces illuminated the mechanisms and predictions of the initial oxygen adsorption step of ORR. They found the relative ORR activity is determined by oxygen access to the platinum surfaces and the slowest reaction steps seem to occur at the same rate independent of the platinum surface. The tools developed in this study can also be applied to other catalyst and electrocatalyst interfaces to enable similar advances in their mechanistic understanding.

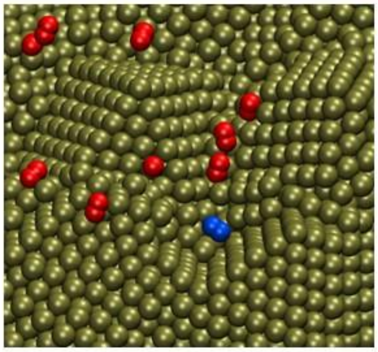

Engineering the atomic-scale surface features of the platinum electrode in contact with the electrolyte helps in attracting oxygen and faster conversion to water. A strongly bound oxygen molecule is highlighted in blue before the reaction. (Image courtesy of Hendrik Heinz and Shiyi Wang, CU Boulder.)

Contact PI: Hendrik Heinz (CU Boulder)
ASCR Allocation PI: Hendrik Heinz (CU Boulder)
ASCR Program/Facility: DD/ALCF
ASCR PM: Saswata Hier-Majumder
Date submitted to ASCR: February 16, 2022
Publication(s) for this work: S. Wang, et. al., *Sci. Adv.* **7**, eabb1435 (2021)

University of Colorado Boulder and University of California, Los Angeles

U.S. DEPARTMENT OF ENERGY | Office of Science

NSF

Binding Energies from Effective Field Theory with Quantified Uncertainties

The Science

A reliable quantitative first-principles description of nuclear structure and reactions with quantified uncertainties remains one of the main challenges in computational nuclear physics. Presently, the most promising approach to reach this ambitious goal comprises a combination of chiral effective field theory (EFT) to describe nuclear interactions in harmony with the symmetries (and their breaking pattern) of QCD with ab-initio few-body methods to tackle the quantum mechanical A-body problem. To address this challenge, the Low Energy Nuclear Physics International Collaboration (LENPIC) aims to develop accurate and precise two- and three-nucleon interactions by pushing the EFT expansion to high chiral orders and using these interactions to solve the structure and reactions of light nuclei.

The Impact

Using up to full-machine runs on Mira and Theta supercomputers, the team performed the first tests of novel chiral nucleon-nucleon potentials with consistent three-nucleon interactions. This demonstrates the importance of three-nucleon interactions and allows for a quantitative understanding of the theoretical uncertainties due to the chiral EFT expansion. The team also extended and tested a Bayesian statistical model that learns from the order-by-order EFT convergence pattern to account for correlated excitations. This allowed them to also demonstrate agreement with experimental ground state energies as well as excitation energies, to within the estimated theoretical uncertainties.

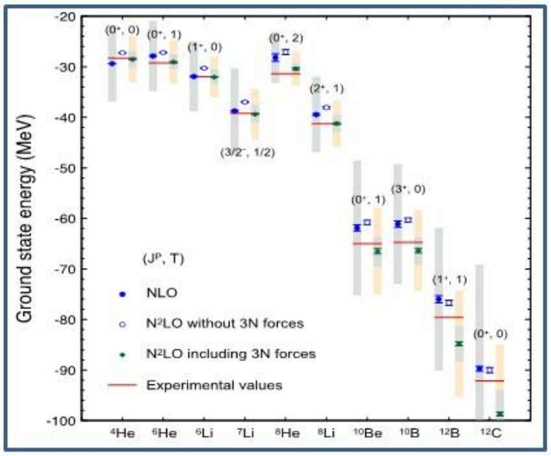

Calculated ground state energies in MeV using chiral interactions in comparison with experimental values. (Image: Evgeny Epelbaum)

Contact PI: Pieter Maris (Iowa State)
ASCR Allocation PI: Gaute Hagen (ORNL)
ASCR Program/Facility: INCITE/ALCF
ASCR PM: Saswata Hier-Majumder
Date submitted to ASCR: March 2022
Publication: P. Maris et al, *Phys. Rev. C* 103, 054001 (2021)

Iowa State University, Ohio State University, Jagiellonian, Ruhr-Universität Bochum, Technische Universität Darmstadt, GSI , Kyushu Institute of Technology, Universität Bonn, Forschungszentrum Jülich, Tbilisi State

U.S. DEPARTMENT OF ENERGY | Office of Science

BMBF, DFG, NSFC, VolkswagenStiftung, PNSC, NSF, DOE

Direct Numerical Simulation of Shockwave/Boundary-Layer Interaction with Crossflow

The Science

High-speed viscous flows are characterized by a thin boundary layer that develops at the point where the flow encounters a solid surface. When a shock (across which the flow velocity and temperature change abruptly), impinges on the boundary-layer, the shock/boundary-layer interaction (SBLI) creates a localized pressure increase which gives rise to an adverse pressure gradient that causes the flow to separate. The separated flow, which may initially be confined to a small separation bubble region, can under certain flow conditions, increase in size and alter the flow path considerably. There is a need to understand the physics of SBLI, in situations that include SBLI interactions with crossflow and flexible surfaces, to develop better models for aerodynamic design of swept wings, control surfaces, and supersonic intakes for jet engines—where SBLI can lead to engine flameout.

The Impact

An important finding from the direct numerical simulation (DNS) of SBLI with and without the imposed crossflow, suggests that the size of the separation bubble can grow by almost 50%. The simulations, which used up to 1,024 nodes of Theta, have also captured the skewing of the flow, as illustrated in the figure (bottom), where due to the presence of cross flow, the local flow angle too can vary quite substantially across the shear layer over the separation bubble. Since SBLI are generally three-dimensional, with some degree of skewed incoming flow, this study holds much promise for developing semi-empirical models for predicting the size of a separation bubble under such conditions.

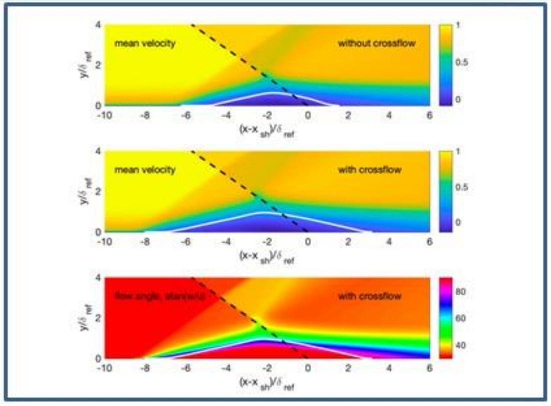

The computed results on top show the shock/boundary-layer interaction without (top) and with (bottom two) crossflow (i.e., out of the page velocity component), with the incident shock marked by the dashed black line. As shown, size of the separation bubble (solid white line) grows by 50% with the addition of the crossflow. (Image: Johan Larsson)

Contact PI: Johan Larsson
ASCR Allocation PI: Johan Larsson
ASCR Program/Facility: INCITE/ALCF
ASCR PM: Saswata Hier-Majumder
Date submitted to ASCR: April 2022
Publication: Di Renzo, M., Oberoi, N., Larsson, J. et al., *Theor. Comput. Fluid Dyn.* (2021).

University of Maryland, College Park

U.S. DEPARTMENT OF ENERGY | Office of Science

Air Force Office of Scientific Research

Extreme-Scale Simulations for Advanced Seismic Ground Motion and Hazard Modeling

The Science

The creation of and dynamics of surface ruptures during earthquakes depend on many factors, such as proximity to fault system, fault geometry, and surrounding material properties (e.g., soil and rock). While simplified models can be helpful to rapidly perform seismic hazard analyses and, for example, inform urban planning and building regulation, empirical fault-displacement models are sparse and poorly constrained due, in part, to the lack of high-resolution observations. Observations from the 1992 Landers earthquake, the third-largest California event of the twentieth century with a magnitude of 7.3, was used to calibrate and validate a dynamic rupture model that reproduced several first-order fault-displacement metrics, such as the location of peak displacement.

The Impact

Fault displacements in large earthquakes have caused significant damage to structures and lifelines and represent a real threat to distributed infrastructure systems spanning faults in more than one location. Physics-based dynamic rupture models and simulations that can capture the general displacement characteristics observed, such as in this work, serve as useful tools to both extrapolate to new scenarios and inform the development of probabilistic fault-displacement hazard analyses. Improved hazard assessments better inform and prepare society for the hazard of earthquakes, enabling mitigation strategies that will have countless societal and economic impacts, ultimately leading to saving lives in the event of a major earthquake.

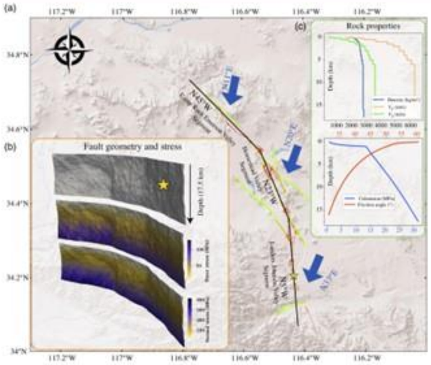

Simulation model setup (a) showing observational and model fault details with arrows indicating orientations of principal stresses, (b) fault-plane model and epicenter (star), and (c) rock properties. (Image courtesy of Yongfei Wang and Christine Goulet, USC)

Contact PI: Yongfei Wang and Christine Goulet (USC)
ASCR Allocation PI: Christine Goulet
ASCR Program/Facility: INCITE/ALCF
ASCR PM: Saswata Hier-Majumder
Date submitted to ASCR: May 9, 2022
Publication: Y. Wang and C. Goulet, *Bull. Seismol. Soc. Am.* **111**, 2574 (2021)

University of Southern California

U.S. DEPARTMENT OF ENERGY | Office of Science

CEC, PG&E, SCEC (NSF & USGS)

Computer Simulations Aided Design of Environmentally Benign Electrolyte for Lithium Battery

The Science

For scientists working to create the next generation of batteries, water has typically been the enemy. If a lithium-ion battery contains any water, it won't work in the conditions needed for it to retain stability and safety. The challenge here is the ability to rationally design an electrolyte immune to water with reduced hazardous risks. A joint experimental and computational approach using molecular dynamics simulations to model the role of water under realistic experimental conditions identified the mechanism by which a novel electrolyte can literally hold water against conventional "dry" electrolytes. The enhanced sampling molecular simulations provided insight into the degree of water clustering, ion distribution, and structure of the cathode-electrolyte interface giving rise to the favorable battery performance.

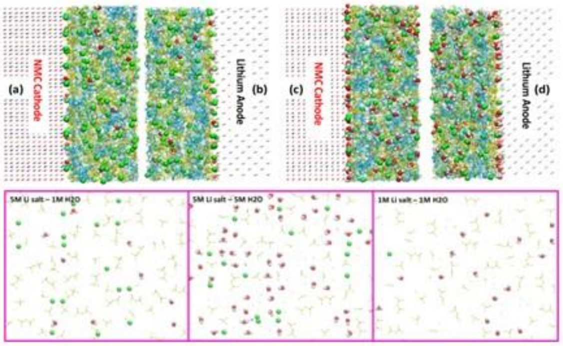

Simulated cluster of water (red) and lithium ions (green) at electrode surfaces under various water concentrations. Battery performance is retained when water molecules remain isolated. (Image courtesy of Wei Jiang, ANL and reproduced from Figures 6 & 7 of *ACS Appl. Mater. Interfaces*, **13**, 58229 (2021).)

The Impact

In state-of-the-art battery manufacturing, critical moisture controls required for electrolytes involve a tremendous endeavor to prepare battery components in a dry environment throughout the process. Electrolyte immunity to moisture in the environment eliminates the stringent control for electrolyte formulation, storage, and transportation. It also facilitates the mass production process and reduces the cost of the battery as the process of removing residual water from battery components consumes a tremendous amount of energy. Researchers demonstrated a novel electrolyte which can accommodate a thousand times as much water as conventional electrolytes and shows exceptional stability on general battery electrodes.

Contact PI: Zhengcheng Zhang (ANL)
ASCR Allocation PI: Wei Jiang (ANL)
ASCR Program/Facility: ALCC/ALCF
ASCR PM: Saswata Hier-Majumder
Date submitted to ASCR: June 6, 2022
Publication: Q. Liu, et. al., *ACS Appl. Mater. Interfaces* **13**, 58229 (2021)

Argonne National Laboratory

U.S. DEPARTMENT OF ENERGY | Office of Science

EERE VTO

Developing a Robust Protocol of Drug Binding Affinity Prediction on Massively Parallel Computing

The Science

The major challenge in drug discovery is the huge chemical space amounting to a very large number of potential ligands which need to be inspected to identify potential drug candidates for each target protein. *In silico* methods play an important role and the most popular are physics-based (PB) and machine learning (ML) techniques, among which the PB techniques are regarded as the foundation of rational drug design and play a central role in database construction/refinement for ML. However, PB methods still have limited applicability in real-world scenarios due to several important yet unresolved issues. A computational medicine group at University College London developed a novel PB approach using molecular dynamics (MD) to realize high-throughput drug binding affinity computations, which was able to provide statistically reliable predictions on massive parallel computers.

The Impact

A high-fidelity high-performance simulation methodology capable of predicting drug binding affinities accurately across huge chemical spaces can substantially reduce the time and cost associated with bringing novel drugs to market. Ensemble-based enhanced sampling MD methods on massively parallel computers are the only way to overcome the chaotic nature of MD and provide reliable uncertainty quantification. Researchers demonstrated the findings from a large dataset spanning broad protein classes and ligand transformations that furbish statistically robust results on the accuracy, precision, and reproducibility. Several physicochemical factors that affect uncertainty were analyzed and definitive recommendations provided concerning the protocols to use that can deliver actionable predictions.

University College London

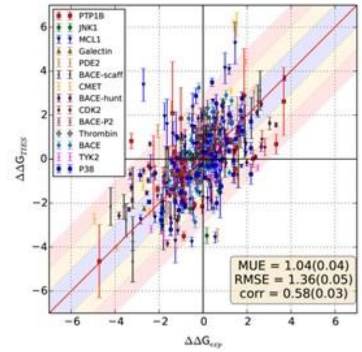

Correlation between experimental and predicting binding affinities for diverse binding complexes, indicated by statistical measures of agreement – mean unsigned error (MUE), root mean squared error (RMSE), and Pearson correlation coefficient (corr). (Image courtesy of Agastya P. Bhati and Peter V. Coveney (UCL) and reproduced from Figure 1 of *J. Chem. Theory Comput.*, **18**, 2687 (2022).)

Contact PI: Peter Coveney (UCL)
ASCR Allocation PI: Peter Coveney (UCL)
ASCR Program/Facility: INCITE/ALCF
ASCR PM: Saswata Hier-Majumder
Date submitted to ASCR: July 19, 2022
Publication: A.P. Bhati and P.V. Coveney, *J. Chem. Theory Comput.* **18**, 2687 (2022)

H2020, NSF, EPSRC, MRC, UCL

Simulating Structural Disordering in Solid-State Battery Electrolytes Using Machine Learning

The Science

Solid-state electrolytes for batteries offer improved safety and increased energy densities; however, atomic disorder at interfaces, grain boundaries, and amorphous regions can severely deteriorate Li-ion conductivity and reduce their effectiveness in real-world applications. Moreover, these disordered interfaces are highly variable and depend on processing conditions, further complicating attempts to probe or control them. Effective computational methods for property prediction must be both accurate and efficient to explore this vast range of possibilities. A machine-learning model was shown to reproduce structural and dynamical characteristics of disordered regions of a solid-state electrolyte, retaining quantum-level accuracy with speeds thousands of times faster than conventional approaches.

The Impact

Only ab initio methods have sufficient accuracy to provide a reliable model for ion transport in complex disordered structures, but the required system sizes, simulation times, and configurational complexity are too large to be practical even with large-scale computing resources. There is a critical need for computational approaches that can accurately simulate a wide range of structure types yet are sufficiently computationally efficient to provide access to dynamical properties in disordered materials. The machine learning (ML) model was able to bridge this gap by accurately representing crystalline, disordered, and amorphous samples of the well-known $Li_7La_3Zr_2O_{12}$ (LLZO) solid electrolyte across a wide range of conditions.

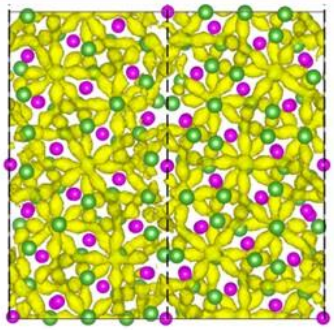

Solid-state diffusion of Li ions through a grain boundary (vertical dashed lines) in the solid-state battery electrolyte LLZO is modeled with a machine learning molecular dynamics simulation. The pathways for dynamic Li are shown in yellow. Green and magenta spheres represent the locations of fixed La and Zr ions in the material. (Image courtesy of Brandon Wood (LLNL) and reproduced from Figure 4 of *J. Chem. Phys.*, **156**, 221101 (2022).)

Contact PI: Brandon Wood
ASCR Allocation PI: Brandon Wood
ASCR Program/Facility: INCITE/ALCF
ASCR PM: Saswata Hier-Majumder
Date submitted to ASCR: August 8, 2022
Publication: K. Kim et. al., *J. Chem. Phys.* **156**, 221101 (2022)

Lawrence Livermore National Laboratory, San Francisco State University, UC Southern California

U.S. DEPARTMENT OF ENERGY | Office of Science

EERE/VTO BMR

Metastable Phase Diagrams for Materials

The Science

Phase diagrams are an invaluable tool for materials synthesis, providing researchers with information on the phases of a material at any given thermodynamic condition (e.g., pressure, temperature, chemical composition). Materials may not reach their equilibrium state during synthesis, operation, or processing and may remain trapped in a local (metastable) free energy minimum. To extend their utility to a promising but mysterious class of materials, researchers from Argonne National Laboratory are using ALCF computing resources and machine learning to develop an automated workflow to construct phase diagrams for metastable materials. The workflow successfully predicted the equilibrium phase for carbon and several observed and predicted metastable structures. The workflow also identified a new metastable phase of carbon that helps to resolve prior experimental observations.

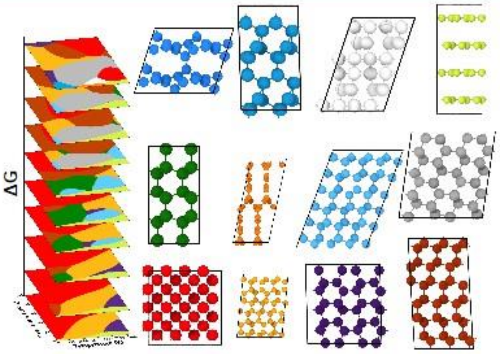

The final product of the machine learning algorithm: metastable phase diagrams for carbon. The colored regions indicated conditions at which carbon exists in certain metastable states (with similarly colored structures) that may yield useful material properties. (Image courtesy of Argonne National Laboratory)

The Impact

Mapping metastable phases and their thermodynamic behavior is highly desirable, but a non-trivial and data-intensive task due to the vast configurational landscape. The team's automated framework for constructing metastable phase diagrams lays the groundwork for computer-aided discovery and design of synthesizable metastable materials, which could help advance a range of applications including semiconductors, catalysts, and solar cells. Metastable phase diagrams not only help to accelerate phase identification by narrowing down the list of potential candidate structures, but more importantly aids in discovering novel polymorphs.

Contact PI: Subramanian Sankaranarayanan (ANL)
ASCR Allocation PI: Subramanian Sankaranarayanan (ANL)
ASCR Program/Facility: DD/ALCF
ASCR PM: Saswata Hier-Majumder
Date submitted to ASCR: September 14, 2022
Publication: S. Srinivasan et. al., *Nat. Commun.* ***13***, 3251 (2022)

Argonne National Laboratory, University of Illinois Chicago, HPSTAR, Northwestern University

U.S. DEPARTMENT OF ENERGY | Office of Science

DOE BES, ANL, AFOSR

Real-Time 3D Analysis During Electron Tomography using *tomviz*

The Science

The demand for high-throughput electron tomography is rapidly increasing in biological and material sciences. This 3D imaging technique is currently bottlenecked by the computational costs of alignment and reconstruction stages which runs from hours to days. The team developed a real-time tomography framework together with dynamic 3D tomographic visualization to enable rapid interpretation of specimen structure as data is collected on an electron microscope. The framework has been integrated into *tomviz*, an open-source and cross-platform 3D data analysis tool that contains intuitive GUI, to enable efficient biological and material structure characterization.

The Impact

The tomviz platform with real-time 3D tomography provides interactive 3D material and biological structure in real-time for enhanced high throughput specimen interpretation. Using geometrically complex chiral nanoparticles, the team demonstrated that volumetric interstation could begin in less than 10 minutes and a high-quality tomogram is available within 30 minutes. In addition, with underlying software support enabling utilizing of parallel GPU-based clusters, real-time tomography significantly improves the throughput of image reconstruction. In addition, real-time tomography allows assessment of the reconstruction convergence on-the-fly and provides visual inspection and intuition of how hyperparameters influence the convergence and the reconstruction quality. Lastly, it also allows early diagnosis of experimental issues such as misalignment, spurious values, and provides guidance for researchers to make adjustment.

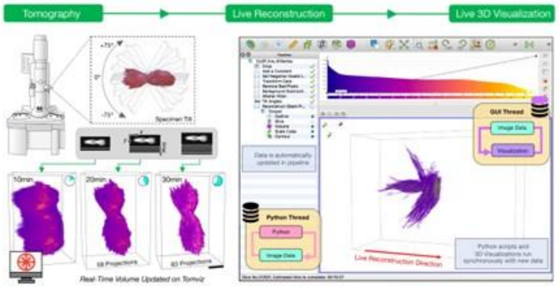

Real-time electron tomography workflow. Specimen projections are sequentially collected in an electron microscope across an angular range (<±75°) and continually passed to *tomviz* for reconstruction and live 3D visualization. As projections accumulate during the experiment, the reconstruction updates in real-time and resolution improves. The *tomviz* platform is composed of a multi-threaded pipeline that synchronously handles tomographic and 3D visualization on separate threads. It monitors for recently acquired data and automatically feeds this into the pipeline. As tomographic reconstructions proceed, visualizations dynamically update and allow for interactive visual analytics.

Figure is reproduced from J. Schwartz et al., Nat Comm. 13, 4458 (2022)

Contact PI: Robert Hovden
ASCR Allocation PI: Robert Hovden
ASCR Program/Facility: DD/ALCF
ASCR PM: Dr. Saswata Hier-Majumder
Date submitted to ASCR: October 12, 2022
Publication: J. Schwartz et al, *Nat Comm.* 13, 4458 (2022)

University of Michigan, Ann Arbor; Kitware Inc; Northwestern University; Cornell University; Argonne National Laboratory; Lawrence Berkeley National Laboratory; Brookhaven National Laboratory

U.S. DEPARTMENT OF ENERGY | Office of Science

DOE, Army Research Office

Direct Numerical Simulation of Flow Over Boeing Speed Bump

The Science

Aircraft design relies on many simulations using simplified models such as delayed detached eddy simulations (DDES). This work identified an important need for re-calibrating those simplified models.

The Impact

Direct numerical simulation (DNS) of flow at Reynolds number $Re=2\times10^6$ over the Boeing Speed Bump highlights an important flow feature not seen at lower Reynolds number, namely that turbulent flows at high-Re do not relaminarize for flows over bump like geometries. The primary reason for the failure of Reynolds-averaged Navier-Stokes (RANS) models to predict separated-reattaching flows rests on the fact that these models were calibrated using DNS results obtained at much lower Reynolds numbers, and on flows lacking the essential physics. Hence, when the old RANS models are used in DDES, the resulting simulations are less accurate at high-Re number. The insights from this DNS are expected to result in new RANS models that, in turn, will make DDES a viable tool in the design of aircraft.

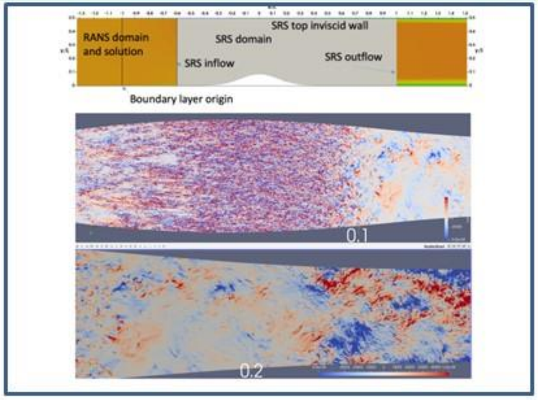

Top: Schematic of computational domain. Middle and lower: Views from above of streamwise vorticity contours on the crest of the bump, and the downstream region where the flow separates. Turbulence in the middle figure shows that at high Reynolds numbers such as this, the flow does not relaminarize at the crest of the bump, unlike at lower Reynolds numbers. (Image courtesy of Kenneth Jansen.)

Contact PI: Kenneth Jansen
ASCR Allocation PI: Kenneth Jansen
ASCR Program/Facility: INCITE/ALCF
ASCR PM: Saswata Hier-Majumder
Date submitted to ASCR: November 2022
Publication: K. Jansen et al., "Direct numerical simulation of the turbulent boundary layer over the Gaussian-shaped Boeing bump at $Re_L=2M$," in *Bulletin of the APS - 2021 Division of Fluid Dynamics Meeting*, **66**, #17, Nov. 2021.

University of Colorado Boulder

U.S. DEPARTMENT OF ENERGY | Office of Science

DOE, NSF, AFOSR, DoD (several grants)

This page intentionally left blank.

Appendix D – Acronyms and Abbreviations

ACA	air channel angle
ACT-SO	Afro-Academic, Cultural, Technological, and Scientific Olympics (NAACP)
AI	artificial intelligence
ALCC	ASCR Leadership Computing Challenge
ALCF	Argonne Leadership Computing Facility
APS	Advanced Photon Source (Argonne)
Argonne	Argonne National Laboratory
ASCR	Advanced Scientific Computing Research
AT	acceptance testing
ATPESC	Argonne Training Program on Extreme-Scale Computing
ATO	Authority To Operate
BIS	Business and Information Services Division (Argonne)
CCIO	custom collective I/O
CELS	Computing, Environment, and Life Sciences (Argonne)
CFD	computational fluid dynamics
CI	continuous integration
COE	Center of Excellence (Intel)
CPS	Computational Science Division (Argonne)
CPU	central processing unit
CSPO	Cyber Security Program Office (Argonne)
CY	calendar year
DAOS	Distributed Asynchronous Object Storage
DD	Director’s Discretionary
DFT	density-functional theory
DHS	U.S. Department of Homeland Security
DL	deep learning
DOE	U.S. Department of Energy
DSAV	Division Site Assist Visit
ECMWF	European Centre for Medium-Range Weather Forecasts
ECP	Exascale Computing Project
EFT	effective field theory
ESH	Environmental, Safety and Health (Argonne)
ESP	Early Science Program
FV&A	Foreign Visits and Assignments (Argonne)
FY	fiscal year
GenSLM	genome-scale language model
GNN	graph neural network
GP	Gaussian process

GPU	graphics processing unit
GUI	graphical user interface
HA	high availability, Hamiltonian Annealing
HACC	Hardware/Hybrid Accelerated Cosmology Code
HDF5	Hierarchical Data Format version 5
HIP	Heterogeneous Interface for Portability
HPC	high-performance computing
HPE	Hewlett Packard Enterprise
HPCM	HPE Performance Cluster Manager
HPSS	high-performance storage system
I/O	input/output
INCITE	Innovative and Novel Computational Impact on Theory and Experiment
ISM	Integrated Safety Management (Argonne)
ISSF	Interim Supercomputing Support Facility (formerly of Argonne)
IT	information technology
JHQ	job hazard questionnaire
JLSE	Joint Laboratory for System Evaluation
JSA	Job Safety Analysis
LCF	Leadership Computing Facility
LDRD	Laboratory-Directed Research and Development (Argonne)
LENPIC	Low Energy Nuclear Physics International Collaboration
LES	large eddy simulation
MD	molecular dynamics
MIT	Massachusetts Institute of Technology
ML	machine learning
MTBI	Mean Time Between Interrupt
MTTF	Mean Time to Failure
MTTI	Mean Time to Interrupt
NDA	nondisclosure agreement
NERSC	National Energy Research Scientific Computing Center
NIST	National Institute of Standards and Technology
NIU	Northern Illinois University
NP	Nonproprietary
NSF	National Science Foundation
Oak Ridge	Oak Ridge National Laboratory
OAR	Operational Assessment Report
OLCF	Oak Ridge Leadership Computing Facility

P3HPC	International Workshop on Performance, Portability and Productivity in HPC
PDU	power distribution unit
PI	principal investigator
PM	preventative maintenance
PMBS	Performance Modeling, Benchmarking and Simulation (SC22)
PNAS	Proceedings of the National Academy of Sciences
PY	previous year
Q&A	question and answer
QCD	Quantum Chromodynamics
QMC	Quantum Monte Carlo
R&D	research and development
RAC	Resource Allocation Council (ALCF)
RDU	Reconfigurable Dataflow Unit
RFI	Request for Information
RMP	risk management plan
RTR	Response to Recommendation
SC	Office of Science (DOE)
SC22	2022 International Conference for High Performance Computing, Networking, Storage and Analysis (annual supercomputing conference)
SDL	Simulation, Data, and Learning
SEC	Simple Event Correlator
SECAC	Secure ASCR Facilities
SME	subject matter expert
SORD	Support Operator Rupture Dynamics
SSL	Secure Sockets Layer
STEM	science, technology, engineering and math
SULI	Science Undergraduate Laboratory Internships (DOE)
TCSB	TCS (Theory and Computing Sciences) Building (Argonne)
UAC	User Advisory Council (ALCF)
UB3	Userbase 3
UC,UChicago	University of Chicago
UIC	University of Illinois at Chicago
WCD	Work Control Documents
WPC	work planning and control
WRF	Weather Research and Forecasting

This page intentionally left blank.

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, Argonne, 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 Rd
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