

Shaping the Future Workforce: Challenges and Lessons Learned in HPC Education from National Labs and Computing Centers

Patrick Diehl
Los Alamos National Laboratory
diehlk@lanl.gov

Ying Wai Li
Los Alamos National Laboratory

Christoph Junghans
Los Alamos National Laboratory

John K. Holmen
Oak Ridge National Laboratory

Elijah MacCarthy
Oak Ridge National Laboratory

Suzanne Parete-Koon
Oak Ridge National Laboratory

Yun (Helen) He
Lawrence Berkeley National
Laboratory

Rebecca Hartman-Baker
Lawrence Berkeley National
Laboratory

Charles Lively
Lawrence Berkeley National
Laboratory

Kevin Gott
Lawrence Berkeley National
Laboratory

Lipi Gupta
Lawrence Berkeley National
Laboratory

Kristina Streu
Argonne National Laboratory

Yasaman Ghadar
Argonne National Laboratory

Paige Kinsley
Argonne National Laboratory

Jane Herriman
Lawrence Livermore National
Laboratory

Erik W. Draeger
Lawrence Livermore National
Laboratory

Victor Eijkhout
Texas Advanced Computing Center

Susan Mehringer
Cornell University Center for
Advanced Computing

ABSTRACT

Workforce training at national laboratories and computing centers is essential and typically falls into two categories: foundational training for newcomers and advanced training for experienced users. Foundational topics—such as version control, build systems, and basic HPC usage—are largely transferable across institutions, while cluster-specific training varies due to differences in hardware, job schedulers, and local workflows. Training on emerging technologies is split between hardware-specific content and broadly applicable programming paradigms. To reduce redundancy and increase impact, national labs, computing centers, and vendors are collaborating through initiatives like the HPC Training Working Group to share best practices, co-develop materials, and broaden outreach. These coordinated efforts aim to make HPC training more accessible, scalable, and consistent across the community.

KEYWORDS

Workforce Development, HPC, Education, Training

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Copyright ©JOCSE, a supported publication of the Shodor Education Foundation Inc.

© 2026 Journal of Computational Science Education
<https://doi.org/10.22369/jocse.2153-4136/17/1/3>

1 INTRODUCTION

The training of the workforce at national laboratories and computing centers has two facets. First, new employees or students, *e.g.*, for internships or summer schools, need to be trained in basic skills. These trainings include build systems, version control, package managers, and how to use high performance compute resources. Second, after users have reached a baseline level of facility with HPC, they need to be prepared for the future with trainings on new hardware, architectures, programming paradigms, and emerging programming languages. Thus, trainings can be categorized into basic skills and emerging technologies. Most basic trainings are common within national laboratories and computing centers. However, some training on how to use compute clusters is different. For example, different job schedulers may be used or some may provide JupyterHub for Python, R, or Julia. Here, most trainings are performed by the local organizations. However, training resources could be shared. For training in emerging technologies, some trainings are specific to a given institution as the institution may have a specific GPU architecture in their supercomputer while other trainings are generic since the technology is architecture independent. Here, vendors or local institutions hold training or hackathons as examples. Additionally, collaborations between institutions for trainings have been established in the last few years. These trainings reflect the current efforts at national laboratories, compute centers, and vendors. Recently, a group of teaching enthusiasts started the HPC training working group¹ and have been

¹<https://olcf.github.io/hpc-training-wg/>

meeting monthly. The goals of the working group are the following: 1) share best practices, experiences, and ideas; 2) share upcoming user training and outreach events; 3) share collaboration opportunities (e.g., on training series); 4) share training materials (e.g., to avoid duplication); and 5) explore platforms for sharing with the broader community, like HPC-ED² [10] and the ACM SIG HPC Education Chapter³.

The paper is structured as follows: an overview of related work is included in Section 2. The background of HPC trainings and current efforts at labs and compute centers are presented in Section 3 and Section 4, respectively. Section 5 describes some multi-center collaborations. Challenges and lessons learned are discussed in Section 6. Section 7 concludes the paper.

2 RELATED WORK

The collaborative group discussed in this paper, consisting of national laboratories and large centers, has been valuable for finding solutions to problems, discussing new ideas, and for building collaborative efforts. There are many other active working groups in HPC training, e.g., the ACM SIGHPC Education Chapter, the CaRCC AI Facilitation Materials Working Group⁴, the ADAC Training, Outreach, & Workforce Development Working Group⁵, the US-RSE Education & Training Working Group⁶, and other groups affiliated with a topic, region, or type of resource. In addition, there are also material and event collections, such as ACCESS-CI Events and Training⁷ and HPC-ED, a metadata catalog of events and materials⁸. Additionally, the Broadening Participation Initiative⁹ [11] of the Exascale Computing Project (ECP) provided internships and mentoring via Sustainable Research Pathways for HPC (SRP-HPC) and enhanced collaboration among institutions via the HPC Workforce Development & Retention Action Group (HPC-WDR).

3 BACKGROUND

3.1 Los Alamos National Laboratory

High-performance computing (HPC) training at Los Alamos National Laboratory (LANL) is conducted under the Institutional Computing (IC) program. A comprehensive set of courses is offered to support users across various levels of experience: 1) Introductory Training: This includes foundational courses on Git, CMake, Spack [4], Unix, HPC fundamentals, and Kokkos [14]; 2) Advanced Training: Topics such as job scheduling with Slurm and strategies for efficient data management are covered; and 3) LANL-Specific Training: These sessions focus on tools and resources tailored to the LANL environment, including Charliecloud [12]; and the use of the open science supercomputer.

Recently, a collaborative training on Introduction to Rust [8] was delivered in partnership with the Pacific Northwest National Laboratory (PNNL). Most training participants are novice HPC

users from a range of divisions. Notably, the Kokkos training spans three days, starting with a beginner-level introduction, followed by two days of more advanced content.

3.2 Oak Ridge National Laboratory

High-performance computing (HPC) training at the Oak Ridge National Laboratory is primarily through the Oak Ridge Leadership Computing Facility (OLCF), which is a Department of Energy (DOE) Office of Science user facility. Training topics largely relate to making use of the facility's exascale Frontier system and include topics such as best practices, debugging, profilers, tools, *etc.* Where possible, the OLCF collaborates with other centers, especially where interests and users overlap (e.g., ALCF, LANL, NERSC). Examples of such collaborations include a performance portability training series with other examples found in the OLCF Training Archive¹⁰. In addition to training events, the facility also hosts hackathons for users to collaborate with staff and vendors on challenges of their choosing, monthly user calls on various topics related to the user experience, and outreach events such as an HPC Crash Course, where participants with little to no background work from Unix basics to hands-on activities on OLCF systems.

3.3 NERSC at Lawrence Berkeley National Laboratory

The National Energy Research Scientific Computing Center (NERSC) is the mission high performance computing facility for the Office of Science in the Department of Energy (DOE SC). NERSC is managed by Lawrence Berkeley National Laboratory and is funded by the DOE SC Advanced Scientific Computing Research Office (ASCR).

NERSC deploys advanced HPC and data systems for more than 11,000 users in 1,000+ projects across a wide range of scientific and computational disciplines. More than 50% of the user base is early career (students / postdoctoral researchers), and 1,000+ new users join NERSC each year.

Training [5] and materials are tailored for a variety of user needs: different user personas, different skill levels, various training categories, and different learning styles; and offered in various styles: presentations, presentations with demos, presentations with hands-on components, hackathons, bootcamps, asynchronous learning with Learning Management Systems (LMS), YouTube recordings with professional captions, short videos, archived training materials (slides, recording, exercises).

NERSC regularly collaborates on training efforts with other HPC centers, especially OLCF, ALCF, and LANL, which has proved highly beneficial to both our staff and users. Collaboration efforts span from basic advertisement to co-developing materials, to co-hosting trainings with the same presentations, and customized hands-on exercises for local systems.

3.4 Lawrence Livermore National Laboratory

HPC training at Lawrence Livermore National Lab (LLNL) is offered by both Livermore Computing (LC) and the HPC Innovation Center (HPCIC).

²<https://hpc-ed.github.io/>

³<https://sighpceducation.acm.org/>

⁴<https://carcc.org/ai-facilitation-materials-working-group/>

⁵<https://adac.ornl.gov/>

⁶https://us-rse.org/wg/education_training/

⁷<https://support.access-ci.org/events>

⁸<https://hpc-ed.github.io/>

⁹<https://www.exascaleproject.org/hpc-workforce/>

¹⁰https://docs.olcf.ornl.gov/training/training_archive.html

As the home of HPC systems and support at LLNL, LC offers trainings that help onboard new LC users and upskill existing LC users. The vast majority of LC trainings are held during the summer, when LC experiences an influx of student interns and thereby new users. For example, “LC Getting Started” is a multi-day workshop introducing HPC, parallel programming, schedulers, GitLab, and Python; shorter workshops target particular internship cohorts. Throughout the year, additional trainings are held as needed to introduce new hardware and tools to all users.

The HPCIC runs an annual online tutorial series featuring LLNL-developed open source projects that span the HPC stack, including performance tools, package management, portability suites, and visualization tools. The HPCIC focuses on outreach, and so this tutorial series targets external (non-LC) users and provides Amazon Web Services (AWS) cloud instances to all attendees. Many of these tutorials serve IDEs in the browser from their cloud instances, such that attendees can log in to a ready environment providing the tutorial stack and materials, simply by navigating to the appropriate web address.

3.5 Argonne National Laboratory

The Argonne Leadership Computing Facility (ALCF) is a U.S. Department of Energy (DOE) Office of Science user facility located at Argonne National Laboratory (ANL). The ALCF provides world-class supercomputing resources and expertise to accelerate scientific breakthroughs across a wide range of disciplines. Supporting thousands of researchers annually through programs such as INCITE¹¹ and ALCC¹², the facility offers access to cutting edge computing systems like Polaris, a pre-exascale system, and Aurora, an exascale system built for the most demanding computational and data-intensive workloads. In parallel, the ALCF invests in workforce development via a year-round training portfolio designed to help new and experienced users make effective use of these systems. Training spans system-specific onboarding and best practices, hands-on workshops and hackathons with staff and vendor experts, and archived materials to support asynchronous learning, with an emphasis on portability, performance engineering, and scalable workflows for Aurora- and Polaris-class architectures. The ALCF also collaborates with peer centers, including OLCF, NERSC, and LANL, to design and deliver joint training that serves overlapping user communities.

3.6 Cornell Center for Advanced Computing

CAC is home to the Cornell Virtual Workshop (CVW)¹³, an online site established in 1995. The CVW includes a wide range of topics relevant to computational research, from programming languages to parallel computing to using large cluster and cloud resources to data science and AI. The materials include quizzes, exercises, a personal notebook, and a large glossary of HPC terms. The CVW has about 250K page views annually. The CVW format has demonstrated sufficient reach and efficacy that it has been incorporated in a number of NSF large-scale infrastructure projects, including

the Leadership Class Computing Facility and Jetstream Research Cloud.

CAC also offers in-person and video conferencing training, primarily for Cornell computational researchers. Lecture materials and shorter tutorial materials are available through our YouTube channel. Most recently, CAC has offered two lecture series per year to all Cornell campuses, attracting about 250 registrations per topic.

3.7 Texas Advanced Computing Center

At the Texas Advanced Computing Center (TACC) at the University of Texas at Austin, HPC training is done by the HPC group. (Other groups at TACC provide other trainings.) HPC trainings range in length from a half day, with topics such as CMake or the Lmod module system, to whole-week HPC ‘institutes’. Since TACC services a US-wide audience, many trainings are online or hybrid; only the institutes are in-person only. In addition to these training classes held at TACC, HPC group members and others also teach for-credit semester-long classes on the UT main campus. The ‘Parallel Computing for Science and Engineers’, ‘Introduction to Scientific Computing’, and several programming classes (targeting C++ and Python) are mainstays of this component.

3.8 Working group

During the 2023 annual Supercomputing conference, staff from three national laboratories met to compare training efforts. These discussions helped expose a variety of commonalities including, e.g., training topics taught, types of events offered, and challenges faced. As a result of these discussions, a monthly meeting was established for HPC centers to share their experiences, practices, ideas, and opportunities for collaboration related to HPC training. Monthly calls have included staff from five national laboratories and three academic HPC centers. Early efforts compared training events across centers to help identify gaps and opportunities to collaborate. Subsequent efforts explored platforms such as HPC-ED and the SIGHPC Education calendar¹⁴ for sharing training material and events. Recent efforts organized cross-center training events. Table 1 summarizes training events and asynchronous materials across centers by category and host institution from 2024 to present. The group has continued to meet monthly and in 2025 formalized themselves as the HPC Training Working Group¹⁵. Monthly meetings are open to all with new centers welcome to the mix.

4 CURRENT EFFORTS AT LABS AND COMPUTE CENTERS

4.1 Los Alamos National Laboratory

In collaboration with ORNL and NERSC, a Julia training session was organized, utilizing Julia notebooks on Perlmutter for hands-on exercises. Building on the success of this initiative, a second iteration is planned for 2026. A series of training sessions was organized in collaboration with NVIDIA. The program comprised four half-day workshops, each dedicated to a distinct topic: Quantum Computing with CUDA-Q [2], HPC SDK, Warp, and Accelerated Python.

¹¹Innovative and Novel Computational Impact on Theory and Experiment

¹²ASCR (Advanced Scientific Computing Research) Leadership Computing Challenge

¹³<https://cvw.cac.cornell.edu/>

¹⁴<https://sighpceducation.acm.org/events/>

¹⁵<https://olcf.github.io/hpc-training-wg/>

Table 1: Training Events and Asynchronous Materials by Category and Host Institution, 2024 to Present. A single training event may cover multiple topics

Training Topics	LANL	LLNL	NERSC	ORNL	TACC	ALCF	CAC
Parallel Programming	3	4	6	11	2	3	8
Debugging	–	2	3	1	1	1	1
Optimization	1	–	2	–	–	3	1
Profiling	1	2	3	2	1	2	1
Visualization	1	2	2	1	–	2	4
C++	1	–	–	1	2	–	1
Fortran	–	–	2	1	–	–	1
Python	1	2	3	7	–	2	6
Performance Portability	3	2	4	4	–	2	–
GPU programming	1	–	3	3	2	2	1
Storage	1	–	2	2	–	2	–
Job schedulers	1	4	–	2	–	3	1
Build systems	1	2	–	–	1	2	–
Module systems	–	–	1	–	–	2	–
Spack	1	2	–	–	–	–	–
Popular OSS	–	4	2	2	–	3	–
Center-specific Resources	1	8	4	2	3	3	6
Workflows	1	2	1	2	1	2	–
Containers	2	–	2	5	1	1	–
Quantum	1	–	3	8	–	–	–
AI, ML, and Data Science	3	2	6	10	–	6	2

4.2 Oak Ridge National Laboratory

The Oak Ridge Leadership Computing Facility’s (OLCF) training team handles training efforts and events at the ORNL. These efforts range from series that run for several weeks to events that run for a few hours. Notable among those are the Frontier Hackathon that spans a week at the minimum where different users of the exascale system participate with their varying science problems. OLCF staff and Frontier COE members join in as mentors and assist the various science teams with their applications with the ultimate goal of optimizing their applications on Frontier. The OLCF also holds new user trainings which aim to equip new INCITE, ALCC and Director’s Discretionary (DD) projects with transitioning and using the centers resources. Periodic series on performance portability, profiling, and GPU programming are also organized, some in collaboration with NERSC and ALCF.

The OLCF also partners with SGCI in organizing hackathons targeting faculties in colleges and universities looking to include HPC content into their curriculum [6]. OLCF also collaborates with Intersect360, AWS, HPE and PSC in holding the Winter Classic Competitions, which introduces students to HPC and provides them with hands-on experience with real HPC systems and applications. Another student-focused training program is the Oak Ridge Cluster Academy (ORCA) which trains participants in HPC system administration.

The OLCF HPC Crash Course, originally developed in 2019 for undergraduate and graduate students with little or no HPC experience, introduces core concepts through modular exercises adaptable to varying coding backgrounds. Foundational programming

content was added in 2020, with AI and quantum computing modules introduced in 2024 and 2025. Delivered 25 times to date in both virtual and in-person formats, the course is actively maintained by the OLCF training team. Materials are also available for asynchronous learning via two GitHub repositories: `foundational_hpc_skills`¹⁶ for exercises that do not require specialized hardware and `hands-on-with-odo`¹⁷ for HPC system-dependent modules.

To engage future talent early, OLCF adapted the HPC Crash Course for high school students in the DOE WDTS-funded Next Generation Pathways program, providing a curriculum in coding, HPC, AI, and robotics to support students’ mentored research projects at ORNL.

4.3 NERSC at Lawrence Berkeley National Laboratory

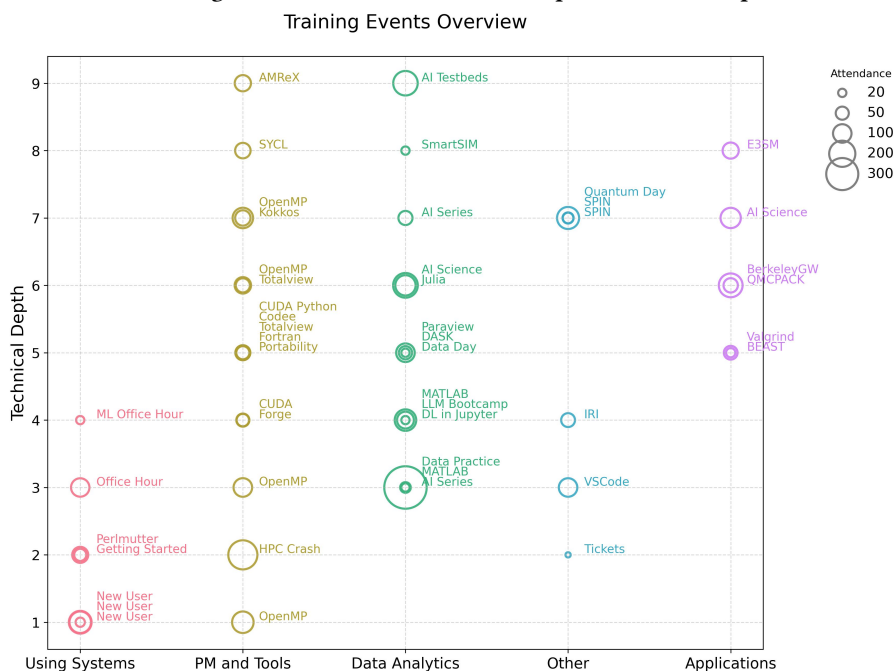
Training is a key component of NERSC’s strategy of providing scalable services to its 11, 000+ users, and NERSC offered a rich set of opportunities for its users to participate in workshops, hackathons, tutorials, outreach, and training events.

Training topics included getting started at NERSC, porting to GPUs, machine learning/deep learning, using performance tools, running applications, programming models, and services. Figure 1 illustrates 2024 training events in five major categories of training topics, organized by level of technical depth and attendance. The above include NERSC hosted/co-hosted training events, and OLCF/ALCF/ECP-hosted training events available to NERSC users.

¹⁶https://github.com/olcf/foundational_hpc_skills/

¹⁷<https://github.com/olcf/hands-on-with-odo>

Figure 1: NERSC Training Events in 2024 with various topics, technical depths, and attendance



4.4 Lawrence Livermore National Laboratory

In recent tutorial improvements, LC has incorporated HPC Carpentry content into its summer tutorials, which has smoothed onboarding for HPC novices.

The HPCIC has been working to improve tutorial deployment in the cloud for our HPC developers. Efforts include wrapper scripts and templates that call the AWS command line interface to ease the provision of AWS cloud instances with minimal configuration or option selection for the developer. We’ve recently created a Slack bot that is configured to launch instances for the target tutorial on demand during a live event; attendees are given access to the bot so that they can provision their own resources, preventing the hosts from over-provisioning resources on the basis of registrations.

4.5 Argonne National Laboratory

The Argonne Leadership Computing Facility (ALCF) develops and curates a wide range of programs and training materials to support high-performance computing (HPC) training, education, and workforce development. Offerings include in-person and virtual events, as well as on-demand resources such as videos, Jupyter notebooks, presentation slides, and example scripts available through Git repositories and the ALCF On-Demand Training website.¹⁸ ALCF programs are designed for users with varied levels of HPC expertise:

- **Introductory training and education:** Covers foundational topics such as parallel programming with MPI and OpenMP, as well as introductions to advanced concepts like using AI for science.

- The Intro to High Performance Computing (HPC) Bootcamp is an introductory program hosted by Argonne and designed for undergraduate students with no prior HPC experience. The bootcamp is a unique multi-center collaboration, and more details can be found in Section 5.2.
- The Intro to AI-driven Science on Supercomputers series is designed to introduce foundational AI concepts to undergraduate and graduate STEM students.¹⁹
- **Intermediate/Advanced training:** Focuses on performance optimization, debugging, profiling, and emerging technologies such as AI/ML workflows on HPC Systems.
 - ATPESC²⁰ is Argonne’s flagship advanced training program. ATPESC is an intensive two-week program for early-career HPC users, featuring a wide range of HPC topics presented by experts from industry, academia, and national laboratories.
 - The ALCF Developer Sessions are monthly webinars connecting ALCF users to developers of leadership-class systems and software.
- **ALCF-specific training:** Provides guidance on leveraging ALCF systems, tools, and services, including system-specific best practices and job scheduling using PBS.
 - The Getting Started on ALCF Systems trainings are live webinars with demonstrations and hands-on experiences for users new to ALCF systems.
 - The INCITE Hackathon is an annual hybrid event for advanced ALCF users associated with INCITE/ALCC projects

¹⁸alcf.anl.gov/support-center/training

¹⁹alcf.anl.gov/alc-f-ai-science-training-series

²⁰Argonne Training Program for Extreme Scale Computing

or prospective applicants. Over the course of 3 weeks selected teams work closely with ALCF mentors and experts from industry to port, optimize, and scale their applications on current HPC architectures.

- The ALCF Hands-on HPC Workshop is an annual user workshop that introduces ALCF HPC resources to the community. This multi day event combines talks, live demonstrations, and guided hands-on sessions to onboard new research teams and accelerate productive use of ALCF resources.

The ALCF also collaborates extensively with other DOE computing facilities including OLCF, NERSC, and LANL, to co-develop and deliver training materials, workshops, and hackathons, detailed in Section 5. These collaborative efforts address the shared needs of the broader HPC community while fostering knowledge exchange among centers.

4.6 Cornell Center for Advanced Computing

CAC is currently working with the Texas Advanced Computing Center (TACC) Leadership-Class Computing Facility (LCCF), contributing to workforce development via developing online training materials for their resource. CAC is also working on online training material for the Chishiki-AI project²¹. The webinar series offered to all Cornell campuses is currently in the planning stages; input from registration forms and exit surveys, along with registration and attendance numbers, are used to inform the topics offered. YouTube how-to videos on using the local cloud resource are currently being updated, driven by a system update.

5 MULTI-CENTER COLLABORATIONS

5.1 Cross-Center Advertising

Training collaborations include events hosted by multiple centers; we also take advantage of training events from other centers to increase training opportunities for our users. Slack channels, Google Spreadsheets, etc. are used to share our events, which can start in the planning stage. For events hosted by multiple centers, each center creates its own event page. For events hosted by one center and open to other centers, the other centers either make mirror event pages or include the events in their weekly announcements to users.

5.2 Intro to HPC Bootcamp

The Introduction to High-Performance Computing (HPC) Bootcamp is a one-week immersive program designed to introduce STEM students to fundamental HPC concepts and their application to solving complex scientific and energy-related challenges.²² Established in 2023 as a collaborative effort among national laboratories, including ANL, LBNL, and ORNL, the bootcamp emphasizes hands-on learning, mentorship, and interdisciplinary collaboration [9]. The bootcamp audience is primarily community college and undergraduate students studying domain sciences with an interest in high-performance computing.

The first iteration of the bootcamp was funded by the Exascale Computing Project, a joint initiative across DOE national laboratories to develop exascale computing ecosystems, with support from the Sustainable Horizons Institute. The 2025 bootcamp was supported and organized by the Argonne Leadership Computing Facility. In collaboration with LBNL and ORNL, national laboratory staff co-developed the curriculum and led hands-on tutorials. Projects focused on real-world HPC and AI applications reflecting ongoing research across DOE mission areas. The program uses inquiry-based activities to demystify HPC and develop skills in parallel processing, cluster computing, and model building. With a mission to inspire the next generation of computational scientists, the bootcamp promotes lifelong learning and creativity and provides a supportive environment for students from all backgrounds to build a sense of belonging and confidence with high-performance computing and artificial intelligence.

5.3 Julia Training Experience

LANL partnered with ORNL and NERSC to deliver a comprehensive Julia training focused on high-performance computing (HPC). The first day featured two talks highlighting how Julia is being applied in production HPC environments at both ORNL and NERSC. On the second day, participants engaged in a half-day hands-on tutorial that included practical exercises.

To support the tutorial, NERSC granted LANL participants temporary access to the Perlmutter supercomputer, enabling them to run interactive Julia workflows directly through JupyterHub.

5.4 Performance Portability Training Series

The increasing diversity of HPC systems has made code portability more desirable. For example, the major HPC systems at ALCF, NERSC, and OLCF feature Intel-, NVIDIA-, and AMD-based GPUs, respectively. Performance portability layers (e.g., Kokkos) have emerged to help address this diversity. Such layers provide developers with portable abstractions that map to underlying programming models, enabling easy transitions across different architectures. To help educate users about portable solutions, ALCF, NERSC, and OLCF hosted a Performance Portability Training Series²³, which ran from 2023 to 2024. The creation of the series was inspired by the overlap in user bases across centers and included both standalone training events (e.g., RAJA, SYCL) and training series (e.g., OpenMP, HIP).

5.5 Open Hackathons

Hackathons pair selected HPC teams with specific technical experts to achieve targeted computational goals during a multi-day event. HPC hackathons have evolved over the years from in-person dungeon sessions to hybrid multi-week events. The Open Hackathon “series” led by the OpenACC organization²⁴, is a popular recurring HPC hackathon series hosted around the world, including at academic institutions and DOE science facilities [1].

The OpenHackathon events have proven to be highly impactful to scientific code teams [3, 7, 13], which is attributable in no small part to their collaborative, open design. With OpenHackathon

²¹<https://www.chishiki-ai.org/>

²²intro-hpc-bootcamp.alcf.anl.gov

²³<https://www.olcf.ornl.gov/performance-portability-training-series/>

²⁴<https://www.openhackathons.org>

events scheduled approximately once a month across North America alone, the OpenACC team is able to adjust the event design, mentor training, application prioritization, and advertising over time to best suit the rapidly changing needs of the HPC community. Each host contributes to this advancement through a review process held after the event, creating a well-recognized, effective event that is available to the HPC community year-round.

Collaborators to these events share advertising for upcoming hackathons, available mentors, computational resources, where appropriate, and discussions to find solutions to a variety of challenges. OpenHackathon participants include ALCF, TACC, NERSC, Pittsburgh Supercomputing Center, NASA, NCSA, NREL/NOAA/NCAR, Georgia Tech University, Princeton University and many others, who have used the shared expertise to achieve a wide variety of organizational goals.

5.6 Cray User Group Birds of a Feather

In 2024, education and engagement team members from Livermore Computing, the Pawsey Supercomputing Research Center, and NERSC led a collaborative session at the Cray User Group meeting, held in Perth, Australia. The session took place during the Programming Environments, Applications, and Documentation (PEAD) Birds of a Feather (BoF), which includes discussion regarding user-facing issues regarding usage of HPC facilities within the HPE/Cray environment. Because an extension of documentation is training and education (which makes documentation accessible), the session focused on actively developing connections between attendees working on similar educational content. By placing physical sticky notes labeled with different popular and relevant HPC topics onto poster papers around the room, attendees could specify which topics they need help creating content for, which topics they have content for, and which topics they need someone else to make and give them content for. Several connections were made during the session by actively engaging the attendees in this process, instead of asking them to make such a list by themselves, on their own time. The information was disseminated out to each participant, as well as contact information for the various attendees and centers, so ongoing collaborations could be fostered in the future.

6 CHALLENGES AND LESSONS LEARNED

6.1 Resources

6.1.1 Community Resources. HPC centers typically maintain center-specific training material archives and announcement mailing lists. For centers, it can be challenging to advertise relevant events from other centers, as including external events may lead to lengthy announcement emails. For users, it can be challenging to navigate several different training material archives when using HPC systems across various centers. Considering the use of community resources for sharing training material and events could be beneficial for easing the user experience. An example for sharing training material could be making use of platforms like HPC-ED [10].

6.1.2 Computing Resources. Some HPC centers have dedicated testbeds for user training. For centers without such systems, it can be challenging to coordinate user training alongside production

user workloads. Similarly, it can be challenging to coordinate user training for new architectures without access to dedicated hardware. Considering collaborating with HPC centers that have access to target hardware could be beneficial for expanding training offerings. Such a collaboration would provide an opportunity to use target hardware before investing in a training system.

6.1.3 Funding Training. Funding requirements for training programs tend to follow a cyclical pattern. In certain fiscal years, significant investment is needed to design and develop new training content. Once the training materials and curriculum are established, funding needs typically decrease, as ongoing costs are limited to delivering the sessions rather than building them from scratch.

A persistent challenge is securing travel funding for external trainers. These trainers – often experts from national laboratories or computing centers – typically volunteer their time and expertise without honoraria. However, they commonly request travel support, which adds to the financial burden.

Another major consideration is the cost of third-party training. While some companies offer limited training free of charge or as part of a maintenance package, many charge substantial fees – particularly for in-person sessions that require multiple instructors on-site. Even virtual trainings can be costly.

One effective strategy for mitigating these expenses is to organize collaborative virtual trainings in partnership with other national labs and compute centers. By sharing both the training and its associated costs, it is easier to justify the expenditure to management and ensure broader access to valuable learning opportunities.

6.1.4 Training Accounts. The combination of attendees from outside the institution participating in a training event and the need to access specific hardware can lead to a problem in handling accounts. TACC previously distributed temporary accounts to attendees who did not already have an account, but this system is now abandoned. Instead, attendees that pre-register are encouraged to create a TACC account, and ‘walk-ins’ can create their account on the spot. For most users, accounts are automatically approved within minutes; someone from user services is at hand for the few cases that need some manual approval, for instance for attendees from certain ‘Countries of Concern’.

At LLNL, LC workshops offer an in-person component where attendees can receive temporary access to existing training accounts and their associated physical tokens for two-factor authentication. This option unfortunately does not work for online participants. LLNL’s HPCIC obviates the need for training accounts on LLNL resources by hosting tutorials on AWS instances, and this approach has allowed us to open attendance to external audiences. Similarly, LLNL employees presenting HPC tutorials at conferences frequently use AWS to provide common computing environments to their attendees.

6.2 Event Planning

6.2.1 Event Format. During the pandemic, virtual meetings became the standard format for HPC training and many other professional events. As some national laboratories continue to support teleworking, offering hybrid training formats has become a logical next step. Interestingly, remote participation remains highly

popular—even among onsite employees, who often prefer to attend sessions from their offices rather than in person. This trend highlights the enduring value and convenience of remote training options in the post-pandemic work environment. While the lecture portions of training sessions are generally effective in both virtual and in-person formats, hands-on exercises present greater challenges, particularly when trying to engage both in-person and remote participants simultaneously. One potential solution is to alternate between fully in-person and fully virtual training sessions, rather than using a hybrid format. This approach could improve interaction and support a more cohesive learning experience for all participants.

When a hybrid training format is chosen, for the consideration of accommodating a large number of in-person attendees (such as events targeting summer students) and still opening the event to remote attendees, one practice we have is always asking in-person attendees to join the online session to fully participate online chat and discussions.

6.2.2 Event Outsourcing. A fraction of our training offerings are outsourced to third parties, such as vendors and developers. Hardware and software vendors of our systems have very close relationships with us. They are experts on various training topics of interest to our users, especially in their software usage and performance optimizations. Such training examples include Nvidia CUDA and OpenACC training series, GPU Programming Bootcamps, and HPE Programming Environment and Tools training. We also have contracts with software vendors that could offer training, such as KitWare (Paraview, Pan3D, CMake), Linaro (DDT, MAP), Totalview, etc. Some organizations have expertise in training topics related to general computing and HPC foundation. One such example is Software Carpentry²⁵ and HPC Carpentry²⁶.

6.2.3 Event Registration. For collaborative training events, we typically mirror event pages across participating institutions to support local promotion. However, managing multiple registration portals complicates logistics. A centralized registration system would streamline organization by simplifying attendee tracking and enabling consistent data collection for reporting purposes. A key challenge is that some national laboratories use internal registration systems that do not permit access to external participants. To address this, a lab-wide registration infrastructure—similar to eduroam in the academic community—would be highly beneficial, allowing seamless identity verification and access across institutions. One interim solution has been to use Google Forms as a centralized registration point. While this simplifies form submission, it lacks robust user authentication, making it difficult to reliably identify and track participants.

The forms always include questions such as “Which HPC centers do you use?”, and record user names for the purpose of keeping track of attendance from each center. It also helps with adding existing users to compute node reservations using a resource allocation for the training event.

²⁵<https://software-carpentry.org/>

²⁶<https://www.hpc-carpentry.org/>

6.3 Materials and Content

6.3.1 Content Difficulty. Pitching cross-collaborative training content to a variety of participant technical competency levels is challenging. To meet learners where they are, training needs to be able to provide integration across multiple modalities, including project-based learning, peer mentoring, asynchronous materials, and scaffolded tutorials, thus creating a flexible and supportive learning environment accessible to participants with diverse technical backgrounds and lived experiences.

In order to accommodate for different skill levels, trainings can be offered at different levels and in multiple parts. NERSC and other centers have adopted a modular or multi-part structure in their trainings—such as CUDA Part 1, Part 2, and Part 3—so that users can engage at the level appropriate for their background without becoming overwhelmed or disengaged. Pedagogical components have been incorporated into trainings to better help users understand what will be learned. For example, NERSC includes Learning Outcomes for its training events. The list of learning outcomes helps users identify the skills that will be learned in each training, enabling the user to determine whether a training is appropriate for them to join. Additionally, layering content delivery across asynchronous materials, guided tutorials, peer mentoring, and project-based activities has proven useful in supporting different learning styles and pacing.

The rapid pace of change in HPC software standards and programming models creates additional challenges in developing content. Updates to the MPI and OpenMP standards (versions 4 and 5), as well as the continuous evolution of CUDA and other GPU libraries, create moving targets for curriculum development. For example, training a user on CUDA Graphs or OpenMP target offload may be irrelevant if their institution lacks compatible compilers and GPUs. This requires site-specific adaptations or disclaimers that add complexity to shared trainings. To address this, some centers like TACC have started offering dedicated sessions specifically on new MPI standards, while others have integrated cross-center collaborations and hackathons to introduce emerging features in a controlled, supported environment.

6.3.2 Hands-On Content. Hands-on activities during training improves knowledge retention and increases engagement. For training events with a hands-on component, we adapt training materials to specific systems. Users can apply the step-by-step procedures learned in the trainings to their own workflows.

The general concepts of hands-on materials are common to all centers in most cases; only the specific user environments of the centers differ, such as default modules, compilers, batch systems, etc. The task of adapting the hands-on exercises to specific centers is usually straightforward when performed by experienced staff from each center.

6.3.3 Performance Portability. Current and emerging HPC systems are becoming increasingly diverse. For centers, it can be challenging to decide which programming model(s) to train users on. For users, it can be challenging to port to new systems when using vendor-specific programming models. The inclusion of performance-portable programming solutions in training offerings could be beneficial for easing transitions between systems. Such inclusion would

help users avoid having to re-write their codes as new systems are introduced.

7 CONCLUSION OUTLOOK

The findings presented underscore the critical role of collaboration between national laboratories and computing centers in shaping the future of HPC workforce development. This collective effort is increasingly important in the context of rapid advancements in artificial intelligence (AI), where HPC serves as a foundational enabler. As AI technologies continue to evolve, the demand for a skilled HPC workforce will grow accordingly, making coordinated initiatives essential for addressing future scientific and technological challenges.

To build on this momentum, the current collaboration will be sustained and strategically expanded. In order to involve additional stakeholders, particularly those national labs and computing centers not yet participating, we plan to organize a Birds of a Feather (BoF) session at either the SC or ISC conference, to serve as a platform to broaden engagement, exchange ideas, and establish new partnerships across the HPC community.

To further enhance the effectiveness of our efforts, we will explore mechanisms for improved organization and sharing of training materials, with the goal of reducing individual workloads and avoiding redundant efforts. Additionally, to strengthen communication with management and program sponsors, we will collaboratively design and implement an effectiveness survey to evaluate and demonstrate the impact of the training initiatives. Let's propel HPC education to infinity, and beyond!

ACKNOWLEDGMENTS

This work was supported by the U.S. Department of Energy through the Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001). Approved by LANL as LA-UR-25-28001 (Rev. 1). This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

REFERENCES

- [1] Izumi Barker, Mozghan Kabiri Chimeh, Kevin Gott, Thomas Papatheodore, and Mary P. Thomas. 2023. Approaching Exascale: Best Practices for Training a Diverse Workforce using Hackathons. *Journal of Computational Science* 14, 1 (2023), 17 – 22.
- [2] Harun Bayraktar, Ali Charara, David Clark, Saul Cohen, Timothy Costa, Yao-Lung L. Fang, et al. 2023. cuQuantum SDK: A high-performance library for accelerating quantum science. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, Vol. 1. IEEE, Bellevue, Washington, USA, 1050–1061.
- [3] Stu Blair, Carl Albing, Alexander Grund, and Andreas Jocksch. 2015. Accelerating an MPI lattice Boltzmann code using OpenACC. In *Proceedings of the second workshop on accelerator programming using directives*. ACM, Austin, Texas, USA, 1–9.
- [4] Todd Gamblin, Matthew LeGendre, Michael R. Collette, Gregory L. Lee, Adam Moody, Bronis R. de Supinski, and Scott Futral. 2015. The Spack package manager: bringing order to HPC software chaos. In *Proceedings of the international conference for high performance computing, networking, storage and analysis*. IEEE, Austin, TX, USA, 1–12.
- [5] Yun He and Rebecca Hartman-Baker. 2022. Best Practices for NERSC Training. *Journal of Computational Science Education* 13, 1 (2022), 23–26. <http://dx.doi.org/10.22369/issn.2153-4136/13/1/4>
- [6] John K Holmen, Je'Aime Powell, Alexander Nolte, Elijah MacCarthy, Charlie Dey, Verónica G Vergara Larrea, Suzanne Parete-Koon, and Linda Hayden. 2024. FacultyHack Events: Faculty-Focused Hackathons for High-Performance Computing Curriculum Development. In *Proceedings of the 8th International Conference on Game Jams, Hackathons and Game Creation Events*. ACM, Copenhagen, Denmark, 67–71.
- [7] Chip Jackson. 2022. Porting OVERFLOW CFD Code to GPUs: To Hackathons and Beyond!. In *OpenACC and Hackathons 2022 Summit*. OpenACC Organization, Virtual, 1 – 20.
- [8] Steve Klabnik and Carol Nichols. 2023. *The Rust programming language*. No Starch Press, San Francisco, California, USA.
- [9] Mary Ann Leung, Katharine Cahill, Rebecca Hartman-Baker, Paige Kinsley, Lois Curfman McInnes, Suzanne Parete-Koon, et al. 2024. Intro to HPC Bootcamp: Engaging New Communities Through Energy Justice Projects. *Journal of Computational Science Education* 15, 1 (02 2024), 49 – 56. <https://doi.org/10.22369/issn.2153-4136/15/1/10>
- [10] Susan Mehringer, Katharine Cahill, Charlie Dey, Brian Guilfoos, David Joiner, Richard Knepper, et al. 2025. HPC-ED: Building a Sustainable Community Driven CyberTraining Catalog. *Journal of Computational Science* 16, 1 (2025), 7 – 13.
- [11] Suzanne Parete-Koon, Mary Ann Leung, Sreeranjani Ramprakash, and Lois Curfman McInnes. 2023. Exascale Computing Project's Broadening Participation Initiative. *The Journal of Computational Science Education* 14 (July 2023), 53–54. Issue 1. <https://doi.org/10.22369/issn.2153-4136/14/1/8>
- [12] Reid Priedhorsky and Tim Randles. 2017. Charliecloud: unprivileged containers for user-defined software stacks in HPC. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC '17)*. Association for Computing Machinery, New York, NY, USA, Article 36, 10 pages. <https://doi.org/10.1145/3126908.3126925>
- [13] Matthew Smith, Arjen Tamerus, and Phil Hasnip. 2022. Portable acceleration of materials modeling software: CASTEP, GPUs, and OpenACC. *Computing in Science & Engineering* 24, 1 (2022), 46–55.
- [14] Christian R. Trott, Damien Lebrun-Grandié, Daniel Arndt, Jan Ciesko, Vinh Dang, Nathan Ellingwood, et al. 2021. Kokkos 3: Programming model extensions for the exascale era. *IEEE Transactions on Parallel and Distributed Systems* 33, 4 (2021), 805–817.