Best Practices for NERSC Training

Yun (Helen) He
NERSC, Lawrence Berkeley National Laboratory
Berkeley, CA
yhe@lbl.gov

Rebecca Hartman-Baker
NERSC, Lawrence Berkeley National Laboratory
Berkeley, CA
rjhartmanbaker@lbl.gov

ABSTRACT

The National Energy Research Supercomputing Center (NERSC) at Lawrence Berkeley National Laboratory (LBNL) organizes approximately 20 training events per year for its 8,000 users from 800 projects, who have varying levels of High Performance Computing (HPC) knowledge and familiarity with NERSC's HPC resources. Due to the novel circumstances of the pandemic, NERSC began transforming our traditional smaller-scale, on-site training events to larger-scale, fully virtual sessions in March 2020. We treated this as an opportunity to try new approaches and improve our training best practices. This paper describes the key practices we have developed since the start of this transformation, including:

- Considerations for organizing events;
- Collaboration with other HPC centers and the DOE ECP Program to increase reach and impact of events;
- Targeted emails to users to increase attendance;
- Efficient management of user accounts for computational resource access;
- Strategies for preventing Zoombombing;
- Streamlining the publication of professional-quality, closed-captioned videos on the NERSC YouTube channel for accessibility;
- Effective communication channels for Q&A;
- Tailoring training contents to NERSC user needs via close collaboration with vendors and presenters;
- Standardized training procedures and publishing of training materials; and
- Considerations for planning HPC training topics.

Most of these practices will be continued after the pandemic as effective norms for training.

KEYWORDS

HPC training, Best practices, Virtual training, Remote training, NERSC, GPU, COVID-19, Zoom, Closed captions

1 INTRODUCTION

The National Energy Research Scientific Computing Center (NERSC) [6] is the primary High Performance Computing (HPC) facility for the Office of Science in the U.S. Department of Energy. NERSC deploys advanced HPC and data systems for more than 8,000 scientists in 800 projects across a wide range of scientific and computational disciplines, including climate modeling, material sciences, fusion, high-energy physics, nuclear physics, biological research, and a host of other scientific endeavors.

The flagship production system at NERSC is the Cray XC Cori system [2], with Intel Xeon Haswell and Intel Xeon Phi KNL compute nodes. The upcoming HPE EX Perlmutter system [10] will contain both AMD CPU and NVidia GPU compute nodes. In preparation for Permultter, we launched the Cori GPU [3] testbed system, a small 18-node cluster of NVidia GPUs, to facilitate porting, benchmarking, and testing efforts for NERSC Exascale Science Applications Program (NESAP) [8] application codes.

NERSC users, including over 1,000 new users annually, possess a wide spectrum of HPC knowledge and familiarity with the usage of HPE/Cray systems. A successful training program for NERSC users is therefore one of the key components of our user support portfolio to enable effective usage of the system resources with efficient programming models. NERSC organizes approximately 20 training events [7] per year.

As a result of the COVID-19 pandemic situation, all of our ongoing training events transitioned into virtual sessions. This gave us an opportunity to try a number of new approaches and iterate on improvements. In this paper, we describe the key practices we have developed during this time.

2 BEST PRACTICES

2.1 Considerations for Organizing Events

Since 2020, NERSC has focused on multi-part training series. This was partly motivated by the constraints of the COVID pandemic; with all events being remote and digital, a format of multiple, shorter sessions was found more effective than full-day or full-week events. The shorter hours per day are more friendly for attendees across different time zones and allow them to have time for their other responsibilities outside the training hours.

For complex topics such as CUDA, OpenACC, or OpenMP, the multi-day format allows users to interleave the consumption of knowledge with the completion of homework. For example, the Deep Learning for Science School in 2020 transitioned from a four-day event to a series of weekly webinars over three months. Similarly, the NESAP hackathons transformed from intensive, single-week events to the same total number of hours spread across a couple of months.

2.2 Collaborating with other HPC Centers and DOE ECP Training Program

Often, NERSC, the Oak Ridge Leadership Facility (OLCF) [9], the Argonne Leadership Computing Facility (ALCF) [1], and the broader Department of Energy (DOE) Exascale Computing Project (ECP) Training Program [5] have common training needs on topics such

https://doi.org/10.22369/issn.2153-4136/13/1/4

as GPU programming and software tools. Collaborating with other training efforts helps increase the reach and impact of events and reduces staff workload for organizing these events individually.

There have been many successful collaborations among these organizations in the past year. A few examples include NERSC users being invited to participate in CUDA and OpenACC training series provided by NVidia and driven by OLCF staff. Meanwhile, OLCF users were invited to join the NVidia HPCSDK OpenMP Offload training and the HPCToolkit training driven by NERSC. NERSC and ECP also collaborated to provide CMake and E4S software stack trainings and invited OLCF and ALCF users. ALCF also included NERSC on some GPU architecture and GPU profiling trainings.

For all these events, staff from the HPC centers and programs worked closely together with the presenters on training logistics, event web pages, compute-system access, and hands-on exercises. Attendees listened to the same presentations and worked on hands-on exercises on various systems to which they had access and on which they already had familiarity with the user environment, with help from HPC center staff and vendors. For participants without a home system or wanting to try something new, we provided NERSC training accounts.

2.3 Targeted Emails to Users

Our training events are announced in the NERSC weekly emails along with many other NERSC center updates. We had found that there was often a tendency for these announcements to be overlooked. Targeted emails to users who may be interested in certain trainings can substantially boost registrations and attendance. We often see a 5–10X registration increase within 30 minutes of sending a targeted email.

For example, new users are contacted for our "Introduction to NERSC Resources" training. Users who have run only single-node jobs according to a Slurm job-accounting analysis are encouraged to take the one-day "Crash Course in Supercomputing" (to learn parallel programming).

2.4 Efficient Management of User Accounts

Management of user accounts for access to special computational resources is often necessary. For example, training attendees may need temporary access to NERSC resources for an event. We have automated the setup of training accounts so a staff member can simply request an event by providing the number of accounts needed and the dates and duration of the event. Then a 4-letter code will be generated for a user to apply for an account associated with the event. The user account will be approved instantly with the login credentials provided. The training accounts will be cleared automatically as well.

Oftentimes, we create compute node reservations for trainings with a hands-on component. It is more convenient for the Slurm batch scheduler for the users eligible for the reservation to be in a small number of projects. In order to satisfy this need, our practice is to add existing NERSC users to the project we use for training, which is allocated from NERSC overhead resources. There are also times when users need to be enabled for a special queue; this is the means of access to our GPU testbed on Cori. At first, the only way to add new users was by hand within the NERSC account management platform IRIS [13], which was quite cumbersome. We have now significantly simplified this process by developing more features such as batch adding or removing users to or from a project or special queue via IRIS APIs.

One major issue with the above approach was that when a user registers for two events that overlap in time, they would have been batch removed after the end of the first event while the user still needs to remain in the training project for the second event. Initially we had to guard these situations manually. After extensive discussions, we have implemented a mechanism that separates individual training events and adds users to separate training projects to prevent any interferences.

2.5 Strategies for Preventing Zoombombing

One of our online events was Zoombombed. After the event, we brainstormed mechanisms to prevent future bombings and incorporated them into our standard practices so that we are prepared for the worst. We recommend the following:

- Use a password on the meeting,
- Do not advertise the URL on the web,
- Enable a waiting room (if possible),
- Turn off participant screen sharing and annotation,
- Mute upon entry,
- Create an unadvertised meeting as backup, and
- Enlist a co-host to help manage attendees.

We also learned that in real time, upon the occasion that Zoombombing happens, there is a "Suspend Participant Activities" button in the "Security" menu that will lock the meeting, disable screen sharing & chat, turn off all audio & video, and give the host a chance to report the incident to Zoom. An excellent resource on this topic, from which we drew many helpful tips, is [12].

2.6 Publishing Recordings with Professional Closed Captions

The Department of Energy is committed to making its information and communications technology accessible to individuals with disabilities. The virtual format for training offers opportunities for improved accessibility, and in 2020 we made it a high priority to make use of these. NERSC made its training materials and recordings public and invested in creating closed captions for the recordings.

By working with NERSC administrative assistants, Berkeley Lab IT and procurement teams, and the support teams for Zoom and Rev (the supplier of captioning services) [11], we have reached a streamlined process to publish professional-quality closed captions for training videos on the NERSC YouTube channel[1]. Fully captioned videos for current training events are already online. We continue to work on retro-captioning our past training videos for a larger impact to the wide user community. These training videos have long-lasting impacts on users and for the public HPC education.

During Zoom meetings, we also enable live subtitles and full transcripts so attendees can toggle them on/off and are also able

[1]The NERSC Training YouTube channel is accessible at https://www.youtube.com/c/NERSCTraining-HPC

to save full transcripts. The real-time captions are helpful; however, they are not as accurate as professional services, and certain custom keywords (such as NERSC) and technical terms (such as OpenMP, Slurm, parallelize, etc.) are often misspelled. These errors are eliminated in the professional-quality captions afterwards with a list of keywords provided by NERSC for the manual captioners.

2.7 Effective Communications for Q&A

A fully virtual event makes Q&A particularly challenging. We adapted to the situation with technology solutions. Slack workspaces were created for most user events and used for Q&A, for sharing presentations, and for continued discussions during and after events. Events not using Slack used a Google Doc in which users could pose questions. When working with external organizations such as HPE and OpenACC, we also used Slack connect to allow people from each workspace to be in the same Slack channel for conversations.

We would like to especially emphasize the multiple benefits of using Slack. We usually create a private #organizers channel for the organizers from different institutions/vendors to use while planning the events. We discuss every detail in this channel before and during the event, planning dates, agenda, system access, presentation slides, hands-on codes, surveys, and more. Using a Slack channel is much lighter and quicker than email exchanges.

Combining Zoom with Slack or Google Docs is more effective for participant communications than Zoom alone. While Zoom does have a chat and Q&A feature, we found these other tools more effective for discussions. As a result, we used Slack or Google Docs for participant interactions in addition to Zoom during the training events. Both are effective for offline discussions and managing Q&A threads. We often had co-hosts monitoring chat and Google Docs and technical experts standing by to answer questions.

Additional practices that improved the user experience included allowing participants to unmute themselves to ask questions in smaller events and during hands-on sessions. We have also used Zoom breakout rooms with screen share and Zoom polls in some events.

At an ECP panel on virtual training best practices [4] held in September 2021, we learned two additional tips. First, in addition to Q&A, Google Docs can also be used for shared note-taking and for quickly gathering attendees feedback such as individual progress for hands-on exercises. Second, Zoom annotation is handy for visualizing feedback; for example, we can ask people to put up a given choice of stamp (green check, red heart, question mark, etc.) that has an agreed-upon meaning.

2.8 Tailoring Training Content for NERSC Users

NERSC put a lot of effort into collaborating with presenters to tailor training content to NERSC user needs (from the largest scope down to small details of materials covered), trying out exercises ahead of time, and providing constructive suggestions. These efforts helped increase the quality of trainings for NERSC users. We also received positive feedback and appreciation from the presenters on developing improved and productive trainings.

For example, the Parallelware training in October 2020 was the culmination of a 10-month collaboration with Appentra staff. The resulting training materials were highly customized to NERSC user needs. We held multiple planning meetings to prepare for this training, studied crucial NESAP user codes and other ECP applications for motifs, provided detailed feedback on training slides and exercises, and performed a survey.

Another example is working with HPE to schedule the standard HPE EX programming environment training at a time close to when users would gain access to the new Perlmutter machine and to collaborate on adjusting the contents of the training to NERSC needs. We established a Slack Connect channel and held multiple long discussions to plan the training, which proved highly beneficial. We agreed to begin with a training for staff focused on concepts that are new with the Perlmutter system (benefiting our system configuration efforts) followed by a short introduction to Perlmutter for users and then an expanded user training with hands-on exercises after users are enabled on the system.

2.9 Standardized Training Procedures

For NERSC HPC consultants and staff from other groups who may provide trainings, we have internally published a guide to hosting a training event that details all the steps and offers guidelines and best practices. It includes detailed information and standard procedures for creating event web pages, registration forms, announcements, and training accounts; details on how to batch add/remove users to/from a project and requesting compute node reservations; and logistics for Zoom, Slack, Q&A, training materials, reminder emails and calendar invitations, welcome and logistics slide templates, Zoombombing prevention, slide publication, production of video recordings with closed captions, post-processing videos, and post-event surveys.

2.10 Considerations for Training Topics

NERSC continues to provide frequent and wide-ranging training opportunities for users. NERSC staff held seventeen distinct training events in 2020 and eighteen events in 2021 as of September. Some were standalone events; others were presented as a series. The courses targeted a broad range of audiences and topics, including getting started with NERSC, GPU architecture, profiling, machine learning/deep learning, tools, running jobs, science applications, programming models (OpenMP, CUDA, OpenACC, etc.), and services. Attendee counts ranged from 10–15 to more than 200 people per event. Feedback from attendees was uniformly positive.

Two courses traditionally offered to Berkeley Lab summer students ("Introduction to NERSC Resources" and "Crash Course for Supercomputing") were opened to all NERSC users in 2021 for the first time, in part because we were no longer constrained by the in-person classroom size limit. Other recent training topics by NERSC included LMOD, CI/CD, checkpointing/restarting, SpinUp, CMake, and OpenMP Offload.

We continue to come up with additional topics on which we may want to offer future training. One recent new training topic was HPCToolkit, and we are beginning conversations with NVidia about adapting for NERSC and offering their Bootcamps, Deep Learning Institute training materials, and Compilers Premier Support. Other topics under consideration are the Perlmutter user environment and

application optimization, Parallelware tools, GPU programming, AI for Science, Kokkos, SYCL, debuggers, etc.

3 CONCLUSION

Remote training has worked better than expected. The virtual training format has proved effective despite the lack of in-person guidance for hands-on exercises and nonverbal feedback from the audience. We were not aware of a good real-time whiteboard sharing tool initially but later found Google Jamboard to be quite useful.

The added advantages of virtual trainings include:

- Larger capacity for training;
- No added administrative overhead of physical meeting planning, site access approvals, and food ordering;
- No user travel needed;
- Ease of arranging multi-session training series;
- Ability to schedule sessions a few days apart without impacting travel
- Shorter training days, so attendees can do some other work during the day;
- Easy to do 1-hour webinars; and
- Increased flexibility for experts, especially for hackathons; some experts can drop in for a partial event, some can present a talk, and some can offer offline Slack help.

After in-person gatherings become safe and commonplace, NERSC will likely still offer more remote training events than in-person ones. We will have some remote-only events and some hybrid events, and can arrange in-person rooms for local people or anyone interested in coming on site, especially for those events with a large hands-on component. We will give remote attendees the same level of attention, if not more than what is given to in-person attendees.

Our recommendation is to continue with most of the practices described in this paper after the pandemic as effective norms of training at NERSC, such as continuing collaborations, continuing shorter and multi-day events, and continuing webinar events.

ACKNOWLEDGMENTS

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.

REFERENCES

[1] 2021. ALCF Training. Retrieved September 29, 2021 from https://www.alcf.anl.gov/support-center/training-overview
[2] 2021. Cori. Retrieved September 2021 from https://docs.nersc.gov/systems/cori/
[3] 2021. Cori-GPU. Retrieved September 2021 from https://docs-dev.nersc.gov/cgpu/
[4] 2021. ECP Panel on Training Virtualization. Retrieved September 29, 2021 from https://www.exascaleproject.org/event/strategies-for-working-remotely-panel-series-training-virtualization/
[5] 2021. ECP Training Events. Retrieved September 29, 2021 from https://www.exascaleproject.org/training-events/
[6] 2021. NERSC. Retrieved September 29, 2021 from http://www.nersc.gov
[7] 2021. NERSC Training Events. Retrieved September 29, 2021 from https://www.nersc.gov/users/training/events/
[8] 2021. NESAP. Retrieved September 29, 2021 from https://www.nersc.gov/research-and-development/nesap/
[9] 2021. OLCF Training. Retrieved September 29, 2021 from https://www.olcf.ornl.gov/for-users/training/
[10] 2021. Perlmutter. Retrieved September 2021 from https://www.nersc.gov/systems/perlmutter/
[11] 2021. Rev: Convert Audio & Video To Text. Retrieved September 29, 2021 from https://www.rev.com
[12] Katie Pratt and Lou Woodley. 2021. CSCCE Tech Tip Sheet - Zoom bombing: How to deal with bad actors during Zoom events. https://doi.org/10.5281/zenodo.4645429
[13] G. Torok, M. R. Day, R. J. Hartman–Baker, and C. Snavely. 2020. Iris: Allocation Banking and Identity and Access Management for the Exascale Era. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–11. https://doi.org/10.1109/SC41405.2020.00046