



# High Performance Computing User Group Meeting - June 12, 2024

June 2024

*Changing the World's Energy Future*

Stephanie J Parker, Ben Nickell, Shane T Grover, Matthew William Anderson, Bradlee Hope Rothwell, Matthew R Sgambati, Brandon S Biggs, Kit J Menlove, Scott M Serr, Trad Ryan Day, Kelly Anne Byrne, Jeremy Alexander Sharapov



#### **DISCLAIMER**

This information was prepared as an account of work sponsored by an agency of the U.S. Government. Neither the U.S. Government nor any agency thereof, nor any of their employees, makes any warranty, expressed 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. References herein to any specific commercial product, process, or service by trade name, trade mark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the U.S. Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the U.S. Government or any agency thereof.

# **High Performance Computing User Group Meeting - June 12, 2024**

**Stephanie J Parker, Ben Nickell, Shane T Grover, Matthew William Anderson,  
Bradlee Hope Rothwell, Matthew R Sgambati, Brandon S Biggs, Kit J Menlove,  
Scott M Serr, Trad Ryan Day, Kelly Anne Byrne, Jeremy Alexander Sharapov**

**June 2024**

**Idaho National Laboratory  
Idaho Falls, Idaho 83415**

**<http://www.inl.gov>**

**Prepared for the  
U.S. Department of Energy  
Under DOE Idaho Operations Office  
Contract DE-AC07-05ID14517**



June 12, 2024

Idaho National Laboratory

# High Performance Computing User Group Meeting



# HPC User Group Meeting Agenda

- Bitterroot Updates and Release to Users
- SLURM
- Open OnDemand Upgrade
- Github Co-Pilot
- Grace-Grace architecture
- HPC Account Renewal System
- Storage Updates
- Sawtooth Updates
- Teton Updates
- Requesting RSICC Codes
- Trainings Supported by HPC

# Bitterroot Status

## Hardware

- 384 Nodes
  - 336 nodes with 256 GB RAM
  - 48 nodes with 128 GB HBM
- 43,008 cores
- 200 Gb/s OmniPath network
- Dell C6620 4-in-2U Chassis as compute
- Dell R760 servers as logins and gateways
- Direct Liquid Cooled on-chip
- Will complement existing systems
- Release date of June 2024

### Commodity Technology Systems-2 (CTS-2)

Deployed at LLNL & Sandia  
Photo credit: Garry McLeod.



# Bitterroot Status Images

- Imaging / Provisioning done with Warewulf4
  - Leverages Apptainer for image building
- Similar to Lemhi
  - Intel
  - OPA fabric
- Slurm
- Memory/Core (256/112 or 128/112)
  - Smaller in-memory image (HBM)
  - Larger image supporting X11 / Desktop

## Commodity Technology Systems-2 (CTS-2)

Deployed at LLNL & Sandia  
Photo credit: Garry McLeod.



# Slurm for Bitterroot

- What is Slurm? (Simple Linux Utility for Resource Management)
  - A job schedule and resource management system similar to PBS
- Slurm is already in use on Hoodoo
- We plan to use Slurm on new clusters going forward
- Many PBS commands (like qstat) will continue to work



# Slurm for Bitterroot

- Submitting a batch job using Slurm

```
sbatch submission_script.slurm
```

- Simple way to start an interactive job

```
srun --nodes=1 --ntasks-per-node=112 --wckey=project_code --time=01:00:00  
--pty -I bash
```

# Slurm for Bitterroot

## PBS

```
#!/bin/bash
#PBS -N hello_world
#PBS -l select=1:ncpus=16
#PBS -l walltime=72:00:00
#PBS -j oe
#PBS -o $PBS_JOBNAME-$PBS_JOBID.log
#PBS -P project_code
```

```
cd $PBS_O_WORKDIR
```

```
module load mpich
mpiexec hello_world
```

## Slurm

```
#!/bin/bash
#SBATCH --job-name "hello_world"
#SBATCH --nodes 1
#SBATCH --ntasks-per-node 16
#SBATCH --time 72:00:00
#SBATCH --output %x-%j.log
#SBATCH --wckey project_code
```

```
cd $SLURM_SUBMIT_DIR
```

```
module load mpich
mpiexec -n 16 hello_world
```

# Slurm for Bitterroot

- Some Slurm commands to be familiar with
  - squeue (shows listing of all currently queued jobs – similar to qstat)
  - scancel (cancels jobs, arrays or steps – similar to qdel)
  - sstat (status information on currently running jobs, including resource use)
  - sacct (tracks and reports job activity and resource usage)
- Stay tuned for upcoming training and documentation

# OnDemand Version Upgrade

- Upgrading from version 1.8 to 3.1
- New features
  - Support tickets
  - One click job submission
  - User profiles
  - Pinned applications
  - Updated files app
- Old functionality should remain



# New Files App

INL Open OnDemand

Files ▾

Jobs ▾

Clusters ▾

Interactive Apps ▾

Information ▾

NCRC ▾

NSUF ▾

Training ▾

>\_ Open in Terminal ▾

Refresh

New File

New Directory

Upload

Download

Copy/Move

Delete

Home Directory

/ home / biggbs / training /

Change directory

Copy path

☐ Show Owner/Mode

☐ Show Dotfiles

Filter:

Showing 1 rows - 0 rows selected

<input type="checkbox"/>	Type ▲	Name	Size	Modified at
<input type="checkbox"/>		exercises	-	2/6/2024 11:19:04 AM

# Built In Support Ticket Form

Linux Desktop (1505131.sawtoothpbs)

1 node | 1 core | Running

Host: sawtooth3

Cancel

Created at: 2024-05-30 12:19:56 MDT

Time Remaining: 335 hours and 59 minutes

Session ID: 6f8fe2df-53b5-4500-8831-ff6cb4eba022

Problems with this session? [Submit support ticket](#)

Compression

0 (low) to 9 (high)

Image Quality

0 (low) to 9 (high)

Launch Linux Desktop

View Only (Share-able Link)

Linux Desktop (2097811.lemhipbs)

Completed

Created at: 2024-05-30 12:22:01 MDT

Delete

Session ID: cc0d0f3e-408e-495c-b00c-c8f510740f19

Problems with this session? [Submit support ticket](#)

For debugging purposes, this card will be retained for 6 more days

Home / Support Ticket

## Support Ticket

Submit a support ticket to INL HPC support staff.

Username

biggbs

Email

Cc

Subject

Attachments

Click the box to attach a file. Screenshots or other files can help troubleshoot your problem.

Add attachment

Description

Submit support ticket

# Interactive Sessions Usage

- Please help us make good use of HPC resources!
- Feel free to leave a "Login"-type desktop session open for its duration.
- Please do not leave "Compute"-type sessions open and unused for significant periods of time.

Interactive Apps

Desktops

Linux Desktop

Linux Desktop with Visualization

GUIs

Barracuda VR

IDE

VSCode Desktop

VSCode Server

Jupyter

Jupyter

RDM

Research Data Management

NCRC

Firefox

Herd URLs

GUIs

NEAMS Workbench

Linux Desktop

version: ab6f8fb

This app will launch a linux desktop on an INL HPC resource.

Project

hpc

This is the project argument provided to the job schedule. For a complete list of projects, go to [projects page](#) on hpcweb

Cluster

Lemhi

Select what cluster you want to run your desktop on.

Job Type

Login

Job Type	Purpose
Login	File editing, compiling, basic testing
Compute CPU	Tasks that may require more than 1 CPU
Compute GPU	Tasks that require GPU resources

Launch

\* The Linux Desktop session data for this session can be accessed under the [data root directory](#).

# Tool to Create Job Submission Scripts

Information

NCRC

NSUF

Training

My Interactive Sessions

Job Options	
Number of nodes:	<input type="text" value="1"/>
Number of cores:	<input type="text" value="1"/>
Number of GPUs:	<input type="text"/>
Walltime:	<input type="text" value="1"/> hours <input type="text" value="00"/> mins <input type="text" value="00"/> secs
Merge output and error files:	<input checked="" type="checkbox"/>
Job name:	<input type="text" value="my_job_name"/>
Project Code:	<input type="text" value="edu_class - University class work and projects"/>
Receive email for job events:	<input checked="" type="checkbox"/> begin <input checked="" type="checkbox"/> end <input type="checkbox"/> abort
Email address:	<input type="text" value="myemail@example.com"/>
queues:	<input type="checkbox"/> router <input type="checkbox"/> short <input type="checkbox"/> gpu

## Job Script

Script format:

```
#!/bin/bash

# Submit this script with: sbatch thefilename

#SBATCH --time=1:00:00 # walltime
#SBATCH --ntasks-per-node=1 # number of processor cores (i.e. tasks)
#SBATCH --nodes=1 # number of nodes
#SBATCH -A edu_class # Project Code
#SBATCH -J "my_job_name" # job name
#SBATCH --mail-user=myemail@example.com # email address

echo "$USER: Please change the --mail-user option to your real email address before submitting. Then remove this line."; exit 1

#SBATCH --mail-type=BEGIN
#SBATCH --mail-type=END

# LOAD MODULES, INSERT CODE, AND RUN YOUR PROGRAMS HERE
```

13

INL/MIS-24-78521

 Nuclear Science  
User Facilities

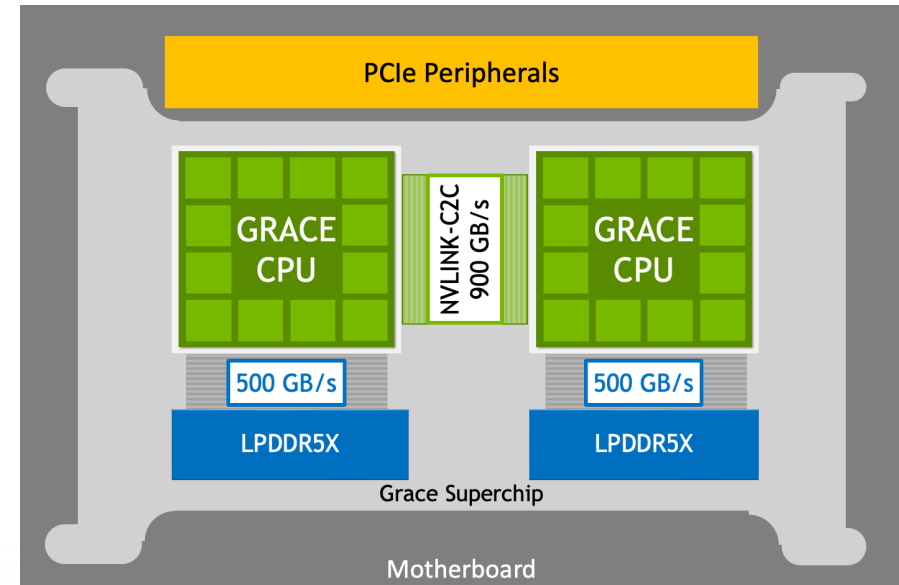



# Roadmap for Local GitHub CoPilot-like Tool

- Monitor feedback and resource utilization
- Support for JetBrains IDEs
- Improved coding models

# Upcoming: Nvidia Grace-Grace

- The Grace-Grace superchip features dual Grace CPUs linked by a high-speed NVLink.
- Grace CPU
  - 960 GBs LPDDR5x memory with ECC
    - 1TB/s bandwidth
  - 144 Arm Neoverse V2 cores
- Total cores 288 and 960 GBs



What you need to know:

- We have acquired two grace-grace nodes, that will be available for testing in the near future. (Q4 2024)
- NOTE: Compilers greater impact performance
  - Upwards of 80% gains depending on the selected compiler

\*product info pulled from Nvidia's publicly published information

# HPC Account Renewal System

- New account renewal process for all external HPC users
  - INL employees can also fill out project report
- New authorship feature

The image shows a screenshot of the HPC Account Renewal System. The top navigation bar includes links for 'Live Status', 'For Users', 'Services', and 'Log Out'. A user profile icon is in the top right. A dropdown menu for 'For Users' is open, showing options: 'User Guide', 'User Dashboard' (highlighted with an orange circle), 'Software', 'Modules', and 'PBS'. Below the navigation bar is the 'DASHBOARD' section with tabs for 'USER INFORMATION' and 'GROUP INFORMATION'. The 'USER INFORMATION' tab is active, showing 'Personal Information' and 'HPC Reports'. The 'HPC Project Report' form is displayed, with fields for 'Report Date', 'Report Name', 'Report Category', 'Report Participants' (with sub-fields for First Name, Middle Initial, Last Name, Organization, Citation Order, and HPC User), 'Scientific Achievement', and 'Significance'. The form includes instructions and a note about public release approval.

## Steps to Renew

Receive an email saying that account is within 30 days of expiring

Go to the HPC Web Dashboard and fill out HPC Project Report

Report will be reviewed by a HPC Admin

Account will be extended for another year

# Isilon Home and Project space



- The Isilon is at **82%** full
  - At the current rate, the Isilon will be full by the end of the year
- The project space is taking **85%**
- Working on freeing up space
- New project space quotas
  - Newly created projects will have a 20TB quota
  - Previous projects will be given a 50TB quota
- Will work with project managers on projects over 50TB



# IBM ESS Scratch



- Fast storage
- Scratch is available on all HPC systems
  - /scratch/<username>
- No quotas
- Files deleted after 90 Days

# Pure Flashblade Software stack



- New S200 Flashblade has been placed into production
- Faster hardware
- Five years support
- Easier upgrade paths

# Sawtooth Update

- Renewed warranty on critical components for one year. (cooling, management, chassis, networking)
- Transition to INL support on compute nodes
  - Proactively filed cases on some potential failures
  - Currently only one of 2079 production compute nodes down
  - 4 tickets still open
- Anticipate additional component failures, primarily memory
  - Harder and more expensive to get DDR4 16GB RAM
  - Trying to procure and test more RAM
  - There are ~25,000 sticks of RAM in Sawtooth compute nodes
- Perhaps 1%-2% of nodes likely to fail and not be repaired in the next year
  - Repairs are time consuming
  - HPC staffing and personnel an issue
  - Time may be better spent making newer hardware available



# Teton

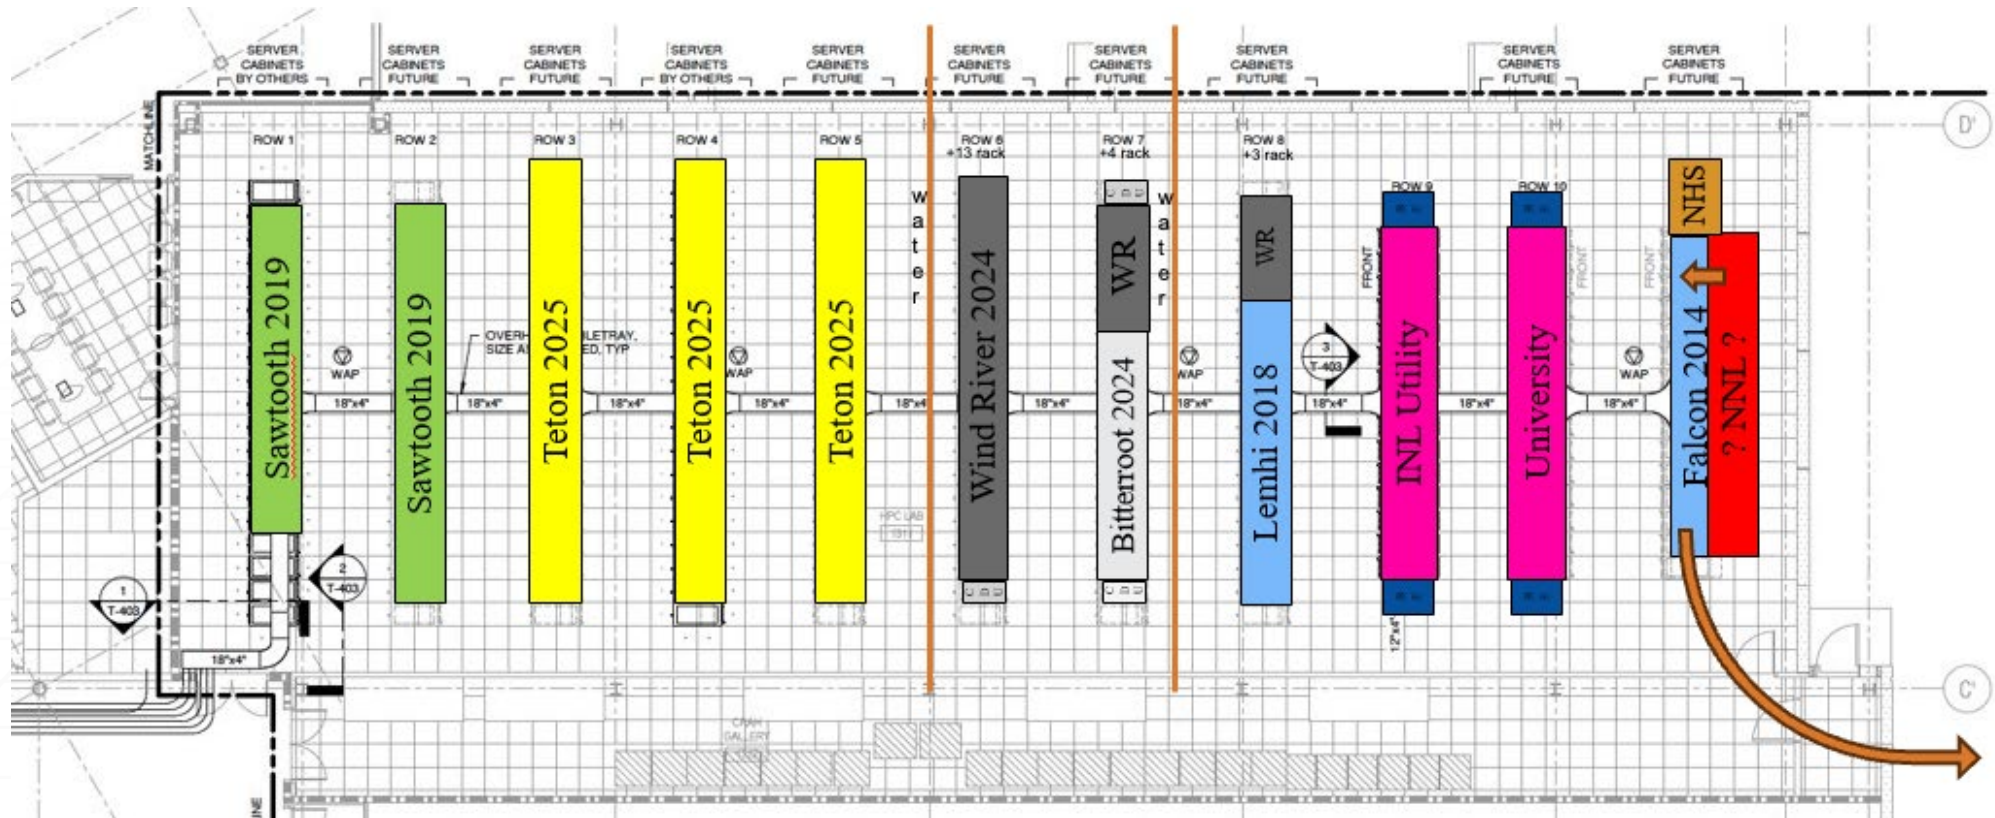
Tetons from the Idaho side  
Photo credit: Ben Nickell

- Next flagship INL HPC system
- Request for Proposals (RFP) is in progress
  - Hope to evaluate bids starting July 2024
- Targeting 400,000 cores at 2GB memory per core
- Fit within 3 rows in INL C3 data center
  - Power: 1.5 MW per row
- Target date: July 1, 2025
- INL HPC and purchasing worked for encourage competition among vendors to get the best possible system
  - Several vendors have indicated that they will be submitting a bid



# Wind River

- Possible INL purchased Dell CTS/2 - (6) Scalable Unit system – arriving late 2024
  - CTS/2 is ~200 nodes and 1.5 petaflops of computing power (per SU)
- Similar hardware to Bitterroot, Intel Sapphire Rapids CPUs, w/o HBM nodes.
- 200GBps Cornelis Networks Omni-Path High Speed Interconnect
- 17.5 racks of compute nodes, direct to chip liquid cooling.



# How to Gain Access to RSICC Codes on INL HPC Systems



- The Radiation Safety Information Computational Center (RSICC) at Oak Ridge National Laboratory is an information analysis center that collects, archives, evaluates, synthesizes and distributes information, data and codes that are used in various nuclear technology applications
- **Licensing:**
  - Register with RSICC and submit a request for code at: <https://rsicc.ornl.gov/>
  - Once a license is obtained from RSICC, go to the Request History link on RSICC's [Customer Service homepage](#)
  - Click on the "**Request History**" link and enter your RSICC Pass Number and Password
  - Once you verify the information is correct, an email will be sent to the address listed in your account
  - Forward this email to INL HPC at [hpcsupport@inl.gov](mailto:hpcsupport@inl.gov), indicating the package(s) to which you want access
  - INL HPC will review and grant access to requested code(s) on INL HPC systems
    - Affiliation on RSICC history must match affiliation for INL HPC account

## Example RSICC History Email:

**From:** pdc@ornl.gov <pdc@ornl.gov>  
**Date:** Wednesday, June 12, 2024  
**To:** User's Email Address  
**Subject:** RSICC User Software/Data History

Our records indicate the following software/data packages have been requested and received from RSICC. If this is incorrect, please contact RSICC at: (865)574-6176 or email pdc@ornl.gov.

Date Requested: 6/12/2024  
Customer ID: xxxxx  
Request Number: xxxxxx  
Name: User's Full Name  
Installation Name: USER'S AFFILIATION, UNITED STATES – This software is only valid for use while associated with this installation.  
Software Package ID: C00834 MNYCP 07  
Software Package Name: SCALE 6.2.4-EXE T

- **IMPORTANT:**
  - RSICC codes are only allowed to be accessed while located in United States. If plans to travel abroad, please contact INL HPC team at [hpcsupport@inl.gov](mailto:hpcsupport@inl.gov) to remove access to codes while on travel.
  - RSICC code is only valid for use while associated with affiliation listed on RSICC history

# INL Nuclear Codes Available for License via the NCRC

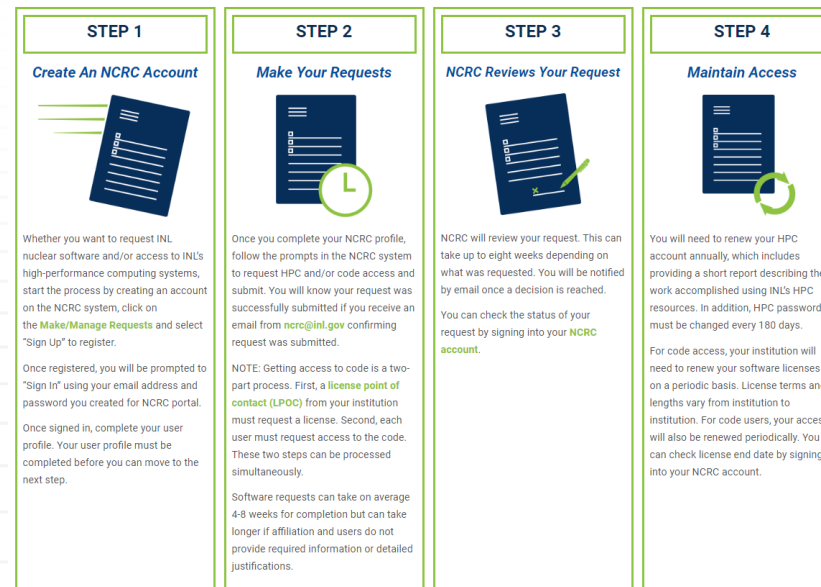
## Available for License

 BISON	>
 BLUECRAB	>
 DIREWOLF	>
 GRIFFIN	>
 GRIZZLY	>
 MARMOT	>
 RELAP5-3D	>
 RELAP-7	>
 SOCKEYE	>
 PRONGHORN	>

## Other Available Codes (Non-INL Licensed)

 SAM	>
---	---

- The Nuclear Computational Resource Center (NCRC) distributes several Nuclear Energy Advanced Modeling and Simulation (NEAMS)-developed applications.
  - 10 INL-licensed software tools require a license by INL to be used by researchers
  - 1 non-INL licensed software tool (SAM) requires a license by Argonne National Laboratory to be used by researchers
- To request access to INL Nuclear [codes](https://inl.gov/ncrc/), go to NCRC portal: <https://inl.gov/ncrc/>



# Trainings Supported by INL HPC

- Availability and access to training on Nuclear Energy Advanced Modeling and Simulation (NEAMS) code is instrumental to continued adoption and refinement of capabilities
- INL HPC staff, in collaboration with NCRC staff, have supported **16** external training sessions and workshops between December 2022 – June 2024 that required HPC access to provide hands on exercises using NEAMS codes for attendees.
  - HPC staff created a HPC account for each attendee, setup the OnDemand app, and scheduled a reservation to meet the requirements for each training
  - HPC staff were made available throughout the training sessions to assist with HPC-related questions from meeting organizers and attendees
- Training opportunities, such as these, promote the development of national computational science expertise

Date:	Training Supported by INL HPC:
December 2022	Bison Training
27-28 June 2023	Bison Training
18 July 2023	Rolls Royce Workshop
16-17 August 2023	Bison Training for X-energy
23-27 October 2023	2023 IAEA Lise Meitner Programme 2 (LMP2)
7-8 November 2023	Bison Training for BWXT
4-8 December 2023	International RELAP5-3D User Group Meeting (IRUG)
11 December 2023	Griffin Training - NRC
18 December 2023	SAM Training - NRC
16-19 January 2024	Bison Fuels Performance Training @MIT January
18 January 2024	SAM Training - NRC
6-7 February 2024	Canadian Nuclear Laboratories (CNL) Bison Training
21 April 2024	ART/NEAMS M&S Pebble Bed Reactors
21 April 2024	ART/NEAMS M&S Molten Salt Reactors
6-8 May 2024	NRC Griffin/Pronghorn Training GC-PBR
16 June 2024	NEAMS Griffin Neutronics Fast Reactor Modeling and Simulation

# Next HPC User Group Meeting

- Wednesday, September 11, 2024
- HPC team will send out dial-in information closer to meeting date



# Questions

- Please type your questions into chat