

17 April 2024

Idaho National Laboratory

INL/MIS-24-77639

High Performance Computing User Experience Updates and Opportunities

HPC: We fight for the users

FIGHT FOR THE USER

- Reduce the barriers to HPC entry
- Provide access to state-of-the-art systems
- Enable massive storage capability
- Give information, resources and support to HPC users
- Aid reproducibility in modeling and simulation
- Deploy machine learning solutions at scale
- Integrate instrument science work into HPC workflows

Open OnDemand

INL OnDemand Files Jobs Clusters Interactive Apps Information NCRC NSUF Training My Interactive Sessions

Help Logged in as andematt Log Out

Notice: If your affiliation has changed since creating your HPC account, please send an update by emailing us at hpcsupport@inl.gov

Quota limit warning for /home/andematt

Reload page to see updated quota. Quotas are updated every 5 minutes.

Using 2.82 TB of quota 3 TB (2.82 TB are yours). Consider deleting or archiving files to free up disk space.

93%



Open OnDemand dashboard is used more than the HPC website!

Open OnDemand

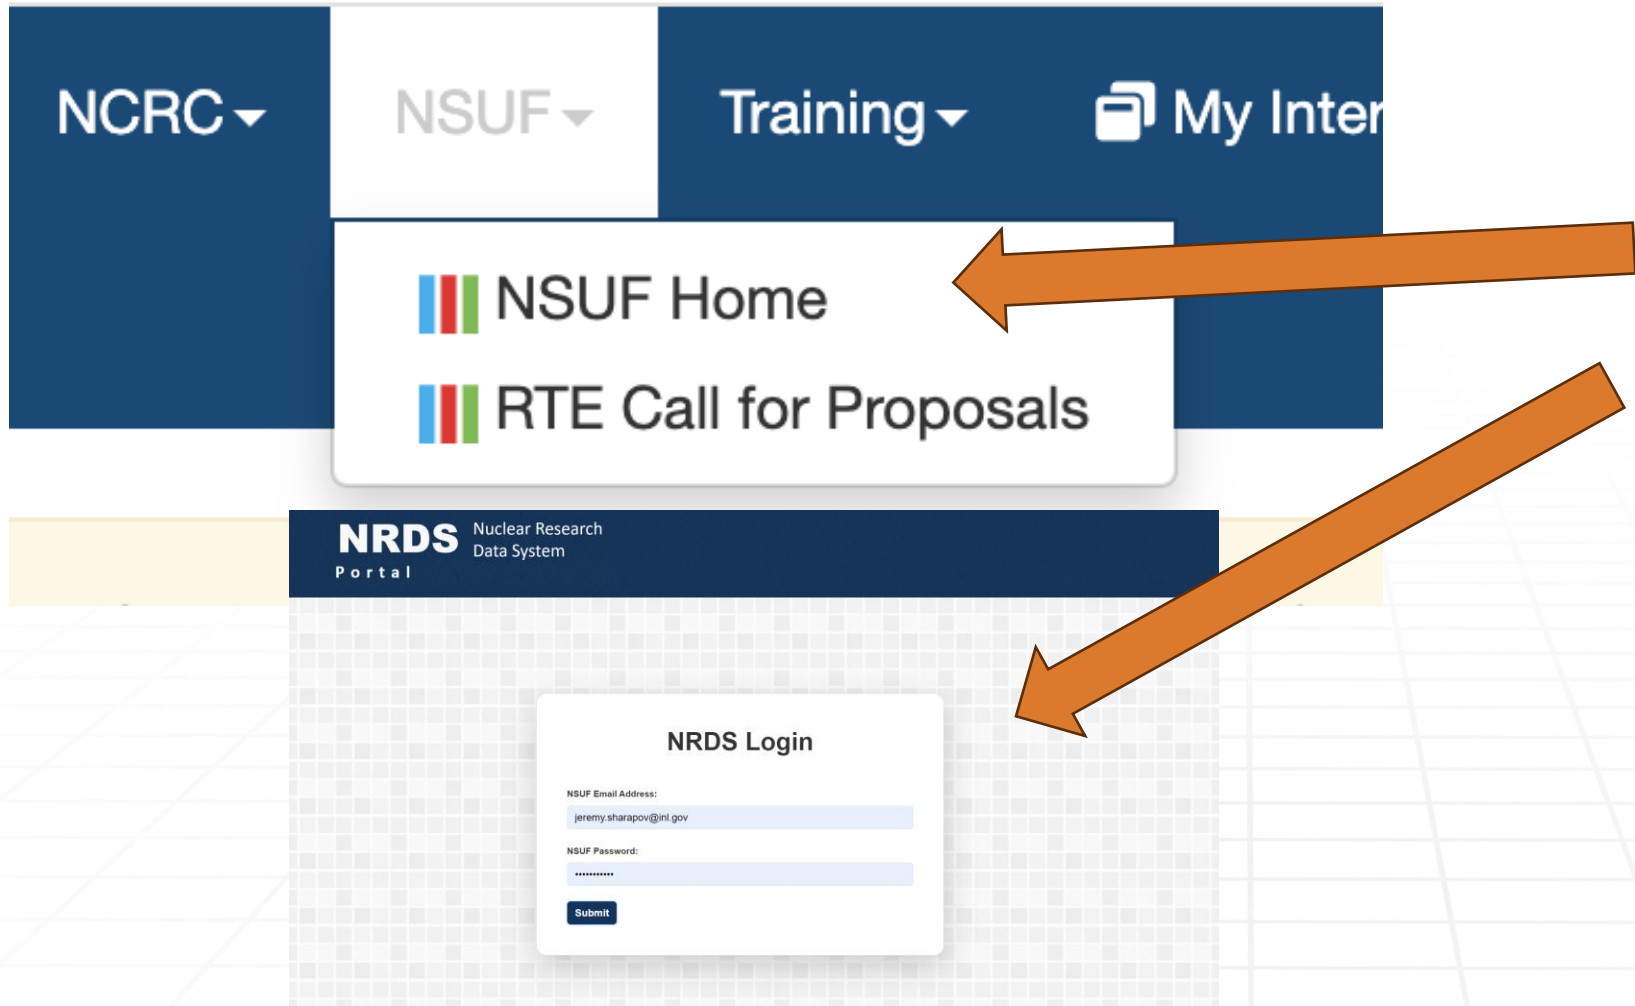
The screenshot shows the Open OnDemand navigation interface. At the top, there are four main menu items: 'NCRC', 'NSUF', 'Training', and 'My Interactive Sessions'. The 'Training' menu is expanded, showing a list of training resources categorized into 'Training Sessions', 'Tutorials', and 'Videos'. An orange arrow points to the 'Canadian Nuclear Laboratory TRISO Training' item in the 'Training Sessions' category.

- NCRC
- NSUF
- Training
 - Training Sessions
 - Canadian Nuclear Laboratory TRISO Training
 - Tutorials
 - AI/ML Symposium 2022 Leaderboard
 - AI/Machine Learning Tutorial Series
 - MIT Symposium Summer 2021
 - Videos
 - AI/ML Training Videos
 - HPC Training Videos
- My Interactive Sessions

Training materials continue to expand

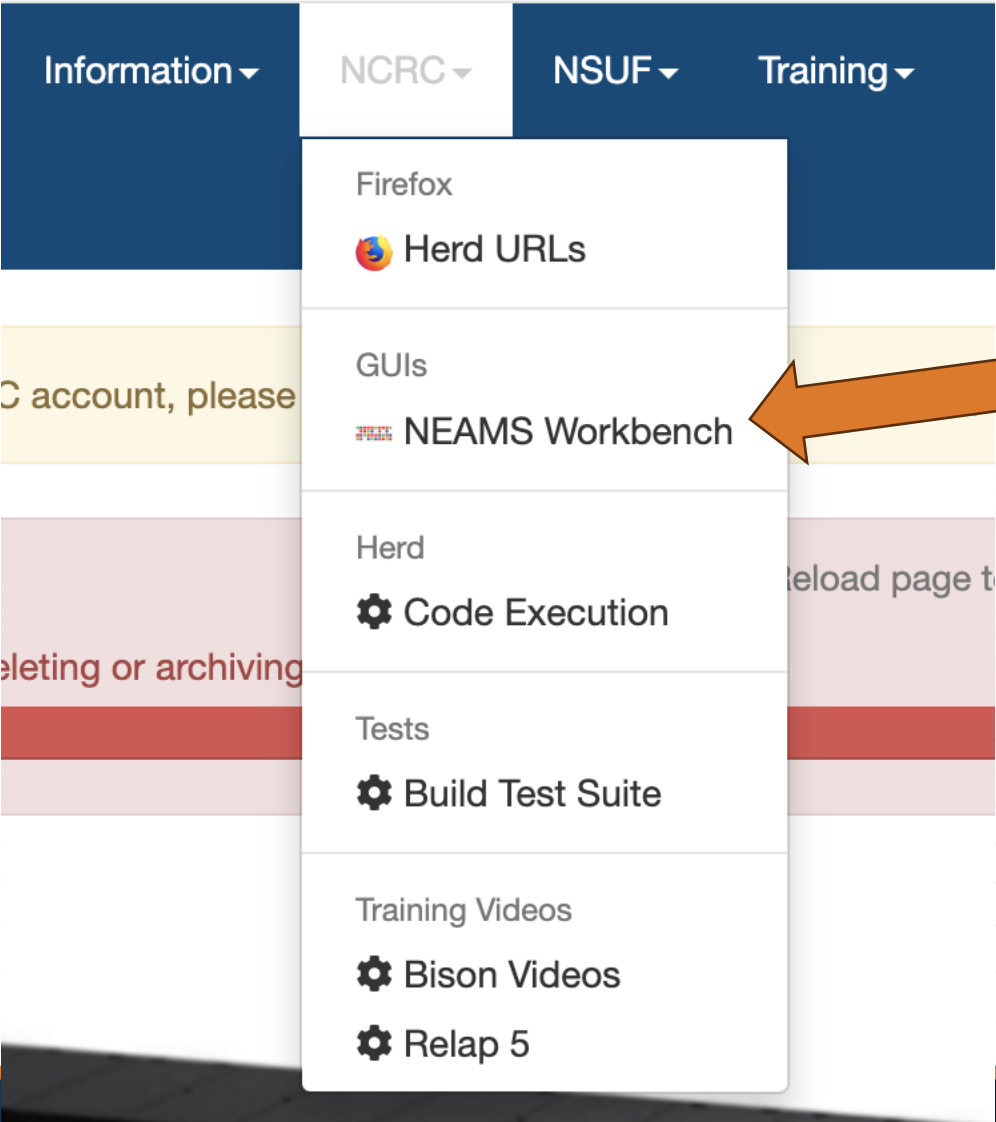
Specific pages for training sessions now standard

Open OnDemand



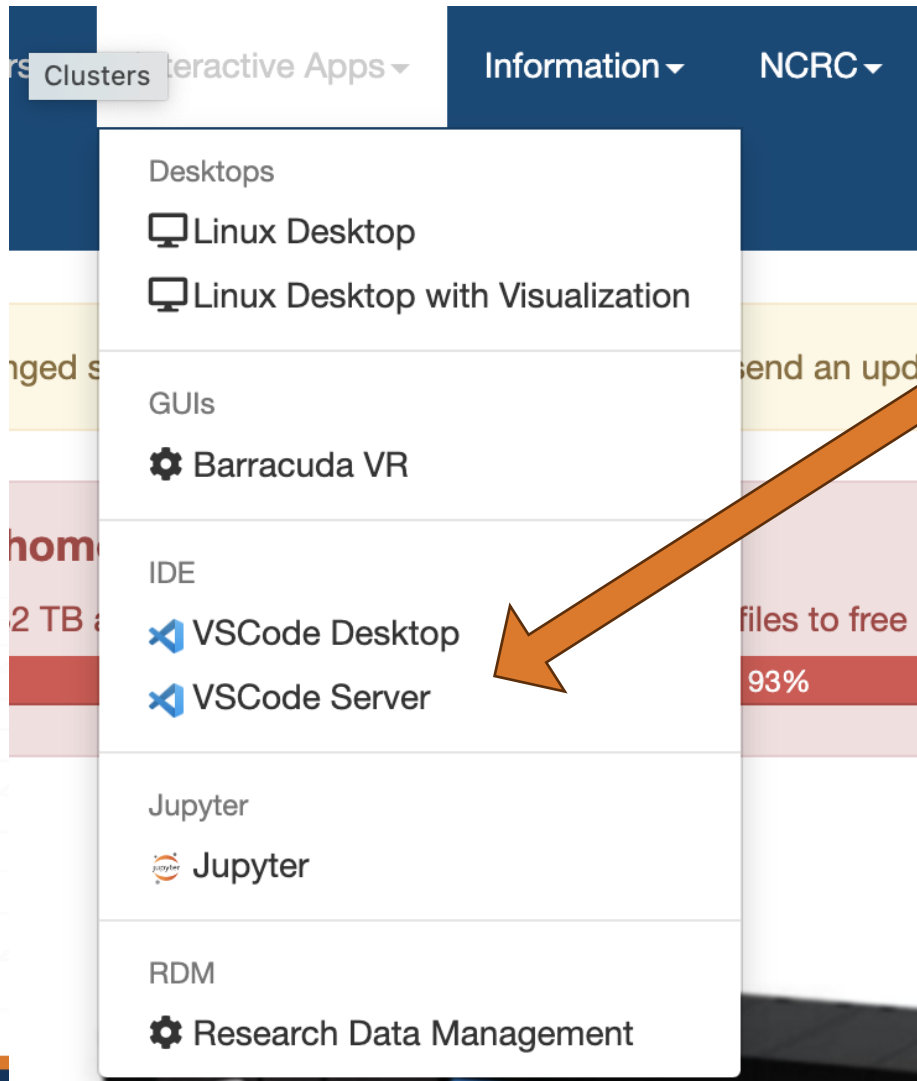
Expanded integration with the NSUF website and call for proposals in both Open OnDemand and the NRDS Portal

Open OnDemand



Expanded integration with the NEAMS training tools, and workbench

Open OnDemand



In-House Continue Extension for VSCode: enables AI code assistant with MOOSE language support

Extension: Continue - Claude, CodeLlama, GPT-4, and more - biggbs - Visual Studio Code

File Edit Selection View Go Run Terminal Help






Search Extensions in Marketplace

INSTALLED

Continue - Claude, CodeLlama, GPT-4, and more - biggbs

Continue

Continue - Claude, CodeLlama, GPT-4, and more

Continue  Continue  Continue  Continue  Continue  Continue

Open-source autopilot for software development - bring the power of C...

Disable Uninstall Switch to Release Version

This extension is enabled globally.

DETAILS FEATURE CONTRIBUTIONS CHANGELOG RUNTIME STATUS

Tab autocomplete now available in pre-release (experimental)

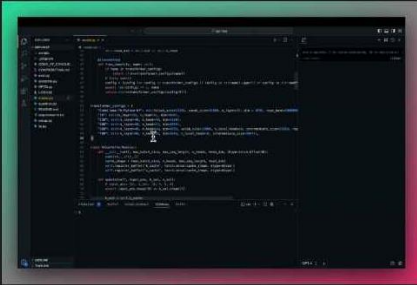
Continue

Continue is an open-source autopilot for VS Code and JetBrains – the easiest way to code with any LLM

License Apache 2.0 continue docs discord join

Features

Ask about highlighted code



Highlight + select code sections and ask a question to get another perspective

- “how can I set up a Prisma schema that cascades deletes?”
- “where in the page should I be making this request to the backend?”
- “how can I communicate between these iframes?”

Reference context inline

Categories

- Programming
- Languages
- Education
- Machine Learning
- Snippets

Extension Resources

- Marketplace
- Issues
- Repository
- License
- Continue

More Info

Published: 2023-05-27, 14:17:59

Last released: 2024-04-11, 13:41:27

Last updated: 2024-04-12, 10:06:12

Identifier: continue.continue

hi

Hello! How can I help you?

Generate a MPI C++ hello world

Here is a simple "Hello, World!" program using MPI in C++:

```
#include <mpi.h>
#include <iostream>

int main(int argc, char* argv[]) {
    // Initialize the MPI environment
    MPI_Init(NULL, NULL);

    // Get the number of processes
    int world_size;
    MPI_Comm_size(MPI_COMM_WORLD, &world_size);

    // Get the rank of the process
    int world_rank;
    MPI_Comm_rank(MPI_COMM_WORLD, &world_rank);

    // Print off a hello world message
    std::cout << "Hello world from process " << world_rank << " of " << world_size << " processes." << endl;

    // Finalize the MPI environment.
    MPI_Finalize();
}
```

To compile this code, you need to have the MPI library installed on your system. You can compile the code with the following command:

```
mpic++ hello_world.cpp -o hello_world
```

To run the code, use:

```
mpiexec -n 4 ./hello_world
```

This will run the program with 4 processes. The `-n` flag specifies the number of processes to run.

Ask a follow-up

+ Add Context Enter

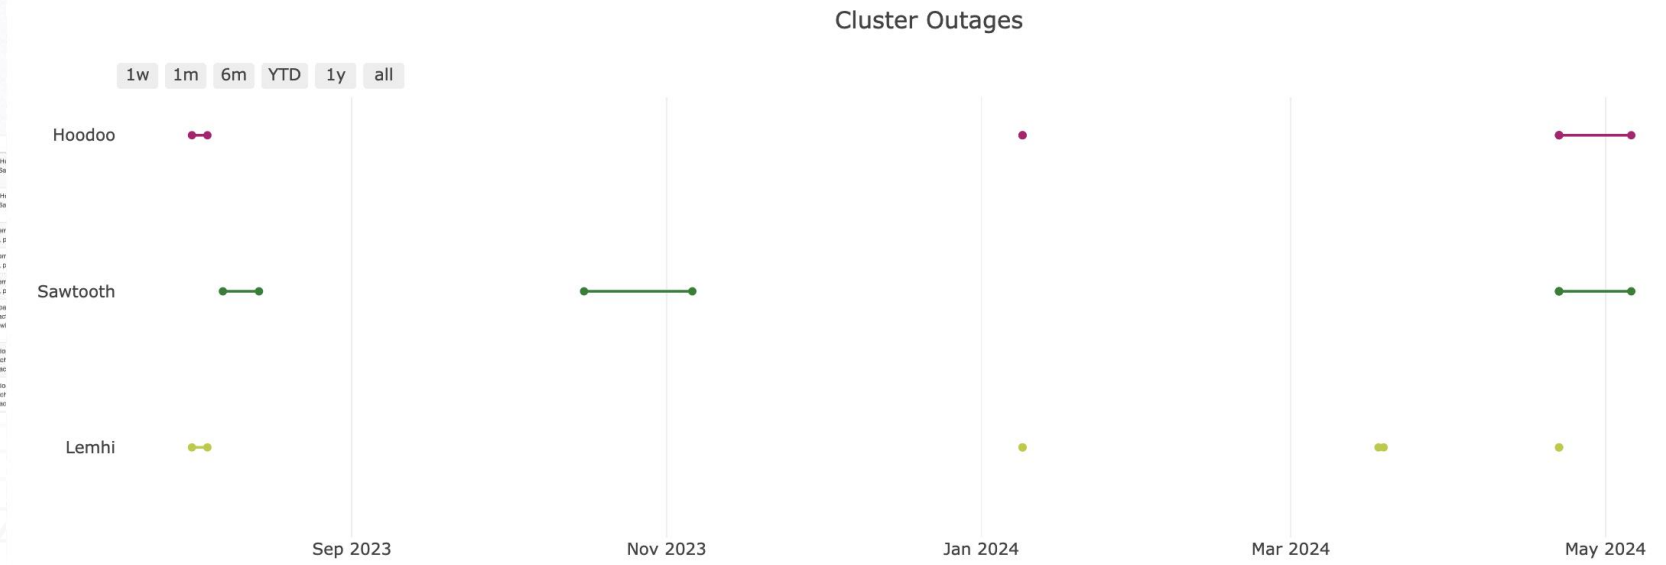
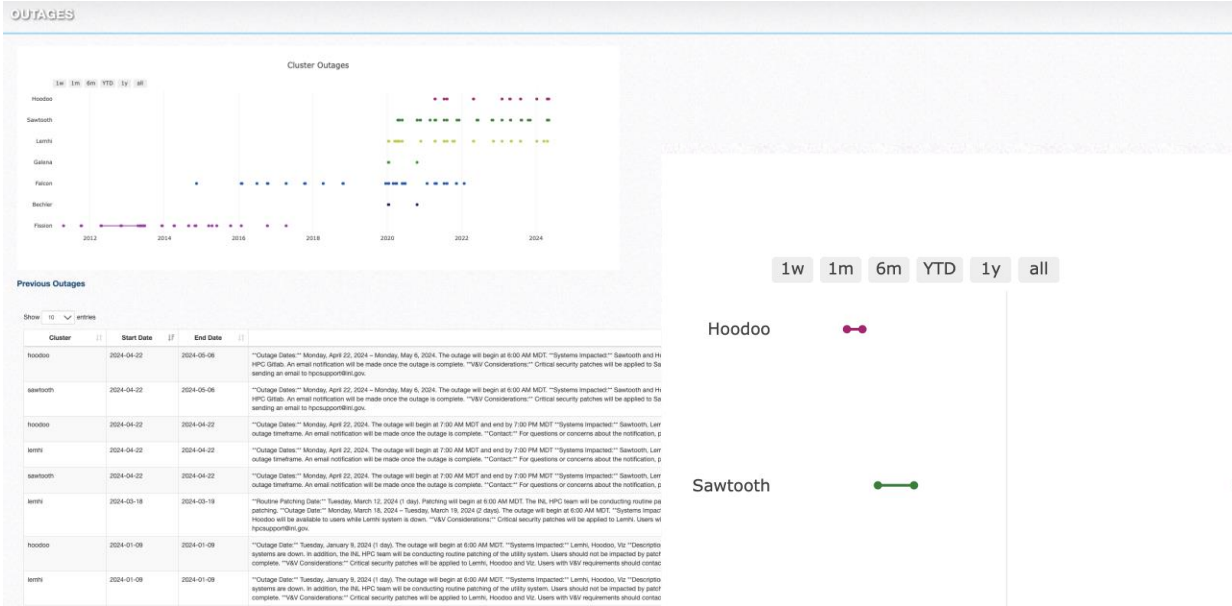
New Session (^ M)

HPC Coding AI +

Continue

Outage Report Page

Tracks all outages, including unscheduled outages



Users are notified of outages 2 weeks in advance or more
Target for unscheduled outages: < 5

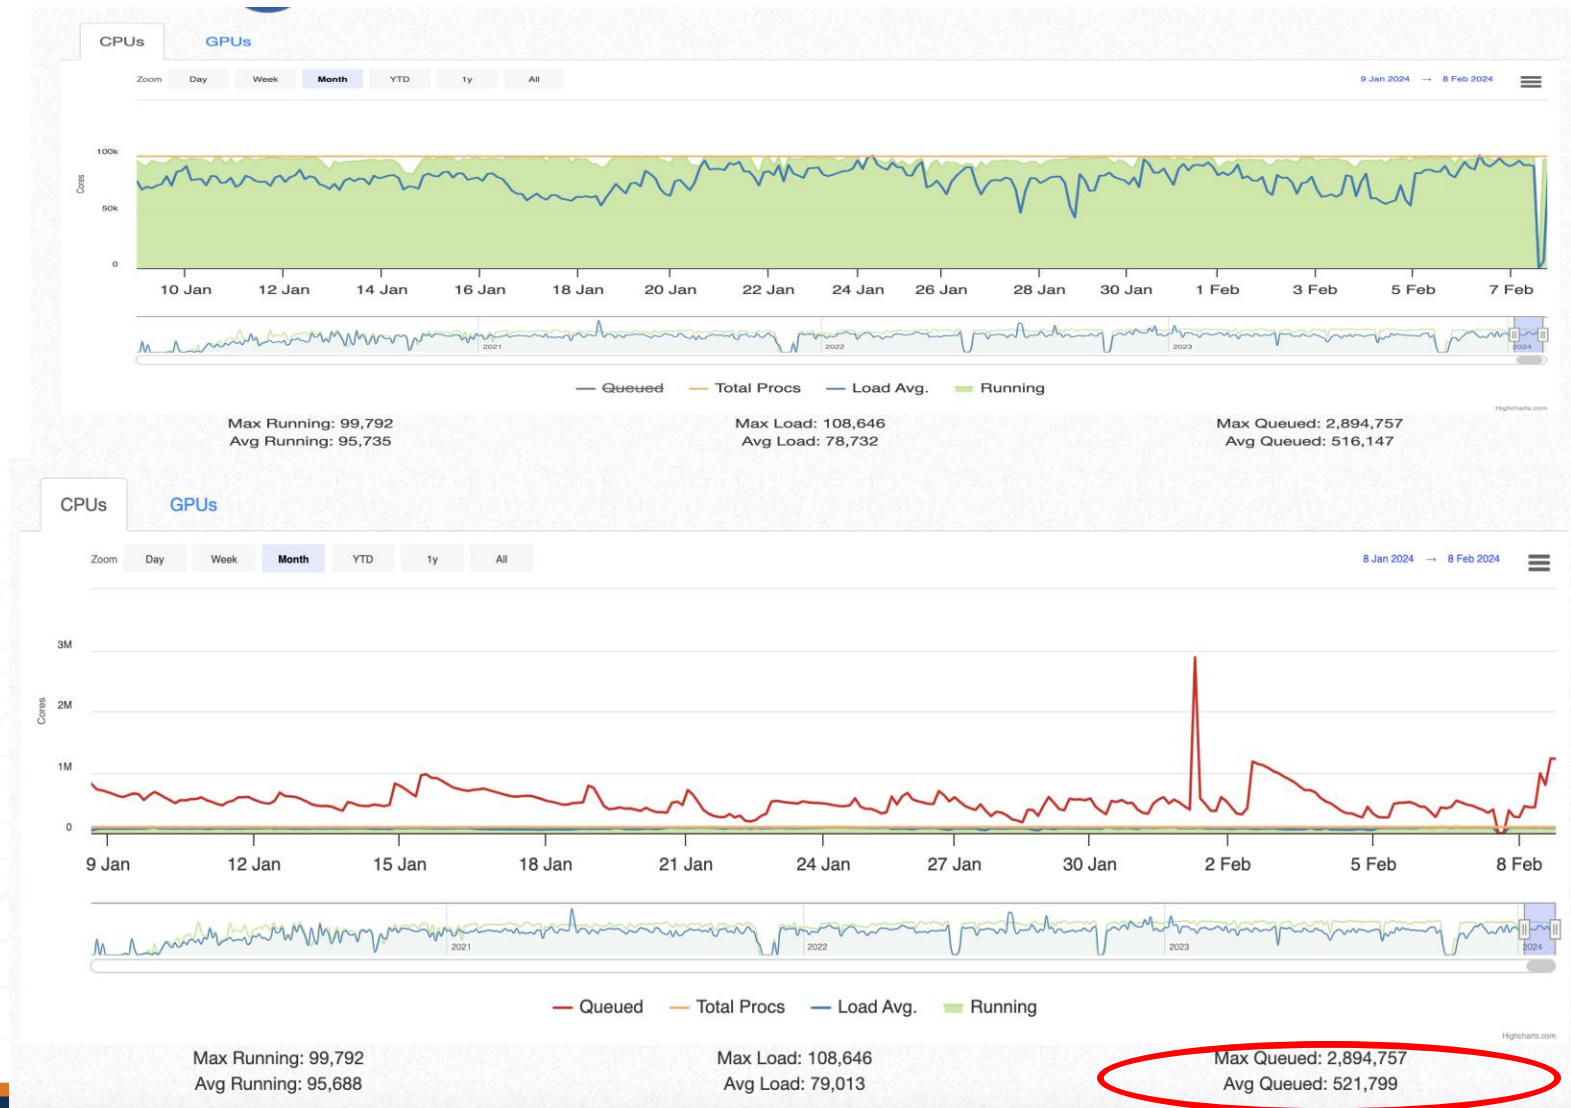
Trainings

FY-24 Trainings using HPC Resources		
10/01/2023 - 09/30/2024		
Date:	Training Name:	Participant Count:
23-27 October 2023	2023 IAEA Lise Meitner Programme 2 (LMP2)	33
7-8 November 2023	Bison Training for BWXT	18
4-8 December 2023	International RELAP5-3D User Group Meeting (IRUG)	84
11 December 2023	Griffin Training - NRC	29
18 Decemebr 2023	SAM Training - NRC	29
16-19 January 2024	Bison Fuels Performance Training @MIT January	31
18 January 2024	SAM Training - NRC	24
6-7 February 2024	Canadian Nuclear Laboratories (CNL) Bison Training	11
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	
16-17 July 2024	SHIFT/NEAMS (SCALE-NEAMS) Training	

Training involves: Reserved queue; project directories; testing; off-hours support; account preparation

Teton Procurement

Current
Sawtooth queue



Explosive growth in demand for certain applications

- Growth in demand for compute resources for NE programs, two examples:

Application	FY22	FY23
Griffin	681 jobs, 48 cores/job	4139 jobs, 97.5 cores/job
MC21	232 jobs, 82.9 cores/job	1016 jobs, 161.4 cores/job

- Growth in demand for AI for NE
 - Hoodoo:
 - FY22: 10.1 million core hours
 - FY23: 15.56 million core hours
- Growth in demand for digital twins simulation support
- Growth in data curation needs (e.g. NRDS)
- Reliable operations

Ticket Response

Requests for support are growing:

- FY21: 4,574
- FY22: 7,019
- FY23: 7,170
- FY24 to date: 3,750

HPC Software Support – Software Installations: Sawtooth

- Installed or Reinstalled Since November 1, 2023
 - Dakota 6.19.0
 - ISOXML
 - OpenFOAM v2312
 - MCNP – occasionally using the gcc has produced segfaults or other errors, trying the intel version may resolve those issues.
 - 6.1.0 - openmpi only
 - 6.2.0 - openmpi/intel both
 - 6.3.0 - openmpi/intel both
 - MFiX
 - SCALE
 - 6.2.3
 - 6.3.4
 - 6.3.1
 - Sierra
 - 5.16.3
 - VASP
 - 5.4.4 with VTST
 - 5.4.4 with Wannier
 - 6.3.0 with VTST
 - 6.3.0 with Wannier

HPC Software Support – Software Installations: Lemhi

- Installed/Reinstalled/Updated Since November 1, 2023
 - Singularity
 - 3.11.5
 - 4.1.1
 - Apptainer 1.2.5
 - Cubit 16.14
 - NBO 7-20210127
 - MCNP6 2.0
 - HDFview 3.3.1
 - go 1.21.6
 - relap53D
 - Perturbo 2.2.0

HPC Software Support – Software Installations: Sawtooth

- Installed or Reinstalled Since November 1, 2023 (cont'd)
 - apptainer/1.2.5
 - boost/1.82.0
 - cuda/12.2.2
 - cutlass/3.4.1
 - deepmd/2.2.7
 - emacs/29.2
 - epic/2020
 - fftw/2.1.5
 - go/1.21.6
 - hdf5/1.14.3
 - jansson/2.14
 - perturbo/2.2.0
 - ImageMagick/7.1.1
 - cp2k/2022
 - 2022-1
 - 2023-1
 - 2023-2
 - 2024-1
 - mpich/3.4.3
 - mpich/4.1.2
 - nccl/2.20.3-1
 - nettle/3.4.1
 - openmc/5feb24
 - pytorch/1.12-geometric
 - relap53D
 - valgrind/3.22.0

Quarterly User Group Meeting

- User group meetings held virtually every 3 months:
 - Dec 12, 2023
 - Mar 13, 2024
 - Next one: June 12, 2024

Avoiding node oversubscription

Goal: Help users use resources efficiently and avoid job failures and resubmissions

- Rule of thumb: Load should not be much > number of cores per node
 - If not: jobs unlikely to run efficiently and/or finish successfully
- Ideal thread count: In an ideal scenario, the number of threads in a process should be equal to the number of cores available. This allows for efficient workload distribution and avoids oversubscription
- OMP_NUM_THREADS
- Every program can be different: If you have questions, reach out!
 - hpcsupport@inl.gov

How do I know if my jobs is oversubscribing?

- Physical status page
- Run 'top' on node
- `ps -eLf | grep <program name>`

<https://hpcweb.hpc.inl.gov/status/sawtooth/physical/>

535.6	50.2	49.8	40.2	48.0	28.3
0.0	33.6	41.5	47.3	43.9	50.4
49.8	49.2	8.0	0.0	44.4	51.1
50.1	49.9	30.9	0.0	49.2	50.2
270.0	32.7	29.3	0.0	49.7	1.0
0.2	28.2	49.7	50.5	32.0	183.7
1199.8	50.4	37.5	0.1	49.7	50.1

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, 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
- **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 to remove access to codes while on travel.
 - RSICC code is only valid for use while associated with affiliation listed on RSICC history



OAK RIDGE
National Laboratory

Sawtooth Maintenance Plan

- Goal: Keep system running as long as possible
- Initial Sawtooth maintenance contract ended February 23rd, 2024
 - Happy 4th anniversary of acceptance and running user jobs, Sawtooth!
- Renewed maintenance support on core components
 - Cluster management software - ensure success of necessary updates
 - Cooling Systems and water chemistry
 - InfiniBand network - high-speed node interconnect
 - Blade Chassis backplane with internal ethernet switches
 - Admin, Login, Viz, Gateways, and Rack Leaders
- What's changing: Compute node maintenance
 - INL HPC staff will do maintenance on compute nodes
 - Buy parts if available
 - Remove - some node attrition as we can't fix them them
 - Repurpose - Combine parts into working nodes to reduce impact



New Pure Flashblade Software stack



- New S200 Flashblade
- Old Flashblade reaching end-of-life
- Faster hardware
- Five years support
- Easier upgrade paths

Questions

