



High Performance Computing User Group Meeting

September 2024

Changing the World's Energy Future

Stephanie J Parker, Ben Nickell, Bradlee Hope Jensen, Brandon S Biggs, Brenyn Faith Rothwell, Devin James Rowley, Kit J Menlove, Matthew R Sgambati, Scott M Serr, Shane T Grover



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

Stephanie J Parker, Ben Nickell, Bradlee Hope Jensen, Brandon S Biggs, Brenyn Faith Rothwell, Devin James Rowley, Kit J Menlove, Matthew R Sgambati, Scott M Serr, Shane T Grover

September 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**

September 11, 2024
Idaho National Laboratory

High Performance Computing User Group Meeting

HPC User Group Meeting Agenda

- FY-24 HPC Annual Report Deadline
- Multiple User Account Policy
- HPC Password Policy
- HPC Geographical Access Policy
- Fairshare and Limits Policy
- Star-CCM Results
- Data Center Outage Dates
- Script Generator Tool
- Lemhi Decommission Date
- HPC Data Center Visual Introduction
- Physical Status Page
- Co-pilot Review
- Training Cluster Update
- OnDemand Review
- Arbiter
- Wind River Update
- Storage Updates
- HPC by the Numbers Graphic Update

HPC Annual Report – User Reports due by October 1, 2024

- The HPC Annual Report highlights users and their projects developed using INL HPC resources.
 - The DOE-NE funding for HPC systems and staff requires users to provide project summaries regarding their work that utilized the INL HPC resources.
 - This requirement is included in the [HPC Access and Usage Policy](#)
 - External HPC account holders are required to renew their accounts every year to continue access to INL’s HPC resources. Internal HPC users are required to submit reports upon request.
- Provide a project summary report of the work you have completed in FY-24 (10/1/2023 - 09/30/2024) using INL HPC resources.
 - All information must be intended for Public Release. This report is made available to the public.
 - Please spell out all acronyms listed in your report
 - All content provided will be formatted by INL staff
- **DEADLINES:** All information is requested to be submitted via the User Dashboard by no later than **Tuesday, October 1, 2024**.
- **EXAMPLES:** See the FY-23 HPC Annual report [here](#), specifically Appendices A – H for user reports submitted for FY-23.



HPC Annual Report – User Reports due by October 1, 2024 cont'd

- **Steps to Renew HPC Account for External Users:**

1. Navigate to <https://hpcweb.hpcondemand.inl.gov/userinfo/dashboard/> in a web browser
2. Enter your HPC username
3. Enter Passcode, which is your HPC RSA SecurID PIN + token code
4. Proceed to enter your HPC username and HPC password
 - If you do not remember HPC password, email hpcsupport@inl.gov
5. Click on the 'Renewal Report' tab
6. Complete ALL fields in the renewal form
 - You will be asked to input publications/conference presentations resulting from your use of HPC resources as well as update user and INL sponsor information, if needed
7. Click 'Add Report'

- **Steps to Submit Reports for Internal Users:**

1. Navigate to <https://hpcweb.hpc.inl.gov/userinfo/dashboard/> in a web browser
 - Must be connected to INL network to access link above
2. Enter your HPC username and HPC password
3. Click on the 'Renewal Report' tab
4. Complete ALL fields in the renewal form
 - You will be asked to input publications/conference presentations resulting from your use of HPC resources as well as update user and INL sponsor information, if needed
5. Click 'Add Report'

The screenshot shows the 'HPC Project Report' dashboard. At the top, there are tabs for 'PERSONAL INFORMATION' and 'GROUP INFORMATION'. Below these, there's a 'Renewal Report' tab. The main section is titled 'HPC Project Report' and contains several input fields: 'Report Name', 'Report Category', and 'Report Participants'. The 'Report Participants' section has a table with columns for 'First Name', 'Middle Initial', 'Last Name', 'Position', 'Organization', 'Division/Order', and 'HPC User'. Below this, there are sections for 'Significance' and 'Key Publications'. At the bottom, there's a 'Report Images' section with a table for 'Image File' and 'Caption'. An 'Add Report' button is located at the bottom right.

Multiple User Account Policy

Multiple Accounts

HPC users with multiple affiliations may, under very limited and controlled circumstances, be permitted to have multiple accounts on INL HPC resources. The intent of this policy is to allow individuals with multiple affiliations the opportunity to manage proprietary data or licensed code access, enabling physical separation of said data or codes. Copying of proprietary data or licensed codes between accounts is strictly forbidden. If other, non-restricted data needs to be copied between accounts, written permission for the originating account sponsor must be obtained in advance.

- INL HPC Appropriate Use Policy: <https://inl.gov/hpc/hpc-policies/>
- A HPC RSA SecurID token is required for each multiple account (external users)
- HPC account renewals & code renewals are required for each multiple account

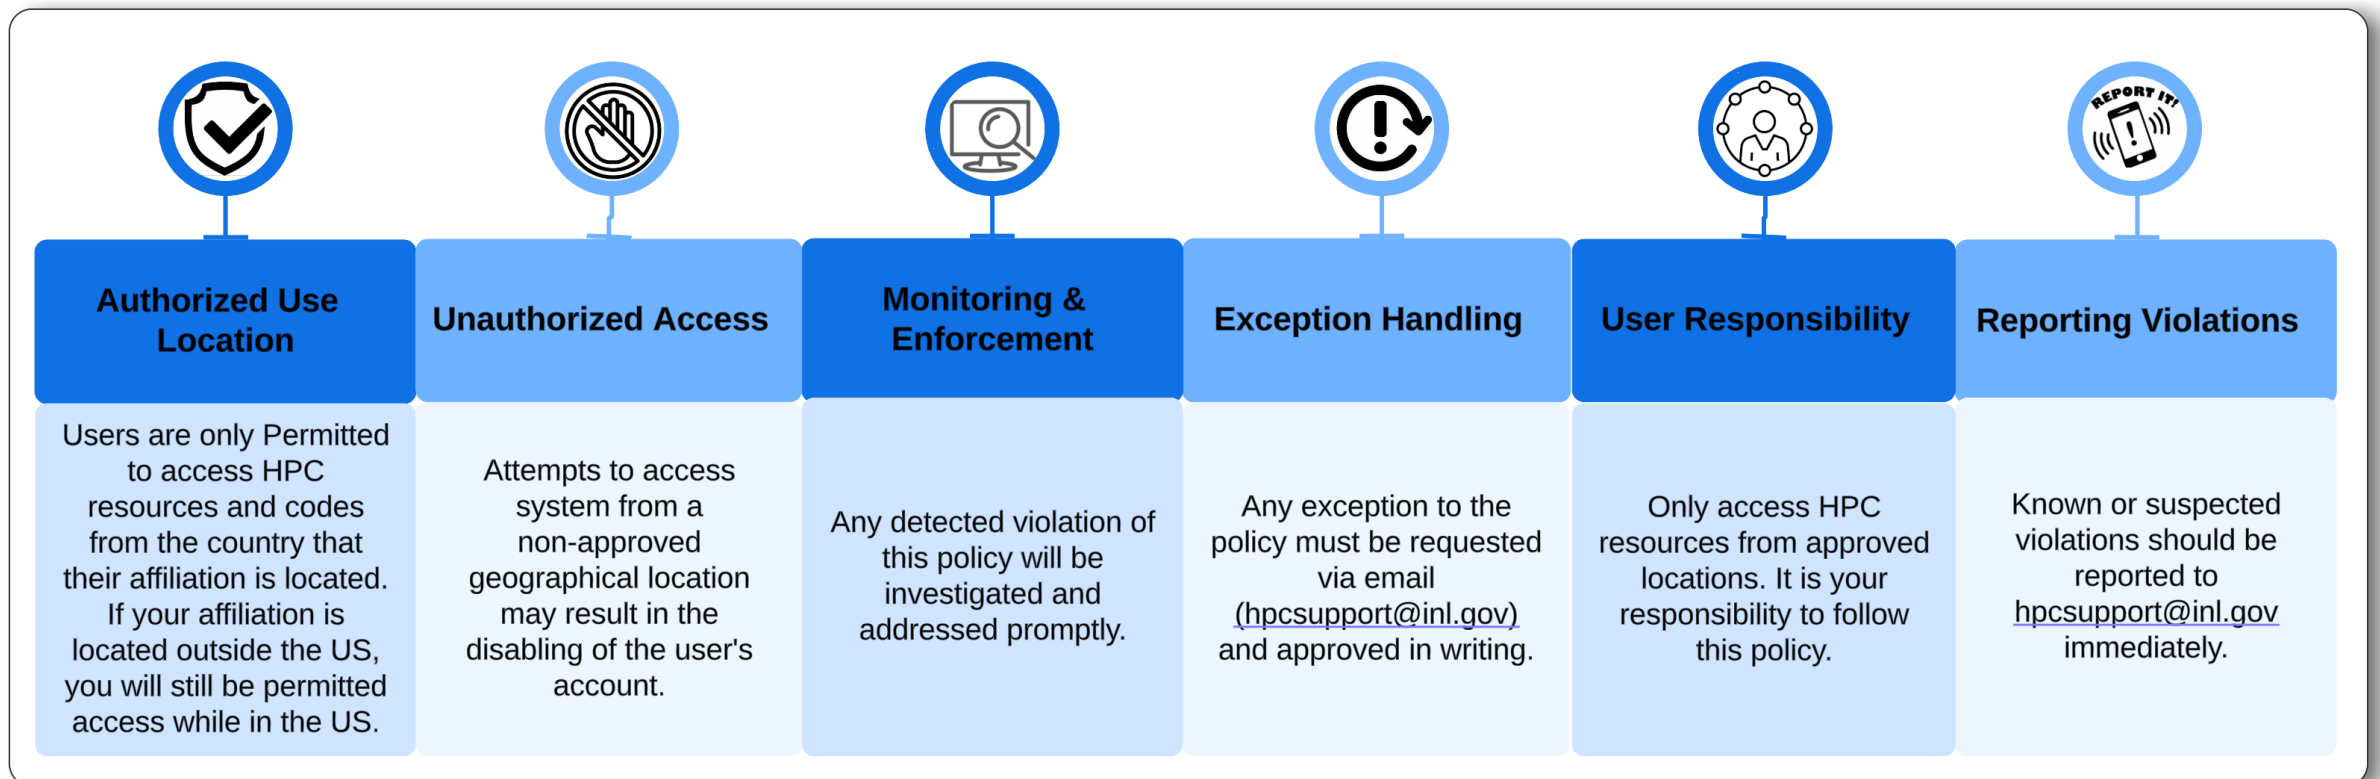
HPC Password Policy Update

- **Purpose:** To update the existing password policy to better safeguard the HPC information assets against unauthorized use, disclosure, alteration, and destruction.
- Change applies to all users of HPC resources (both internal and external)

Major changes	Content	Complexity	History	Tips
<ul style="list-style-type: none">~Minimum of 16 characters~Passwords have no set expiration timeframe~Passwords should not be the same as your INL account	<ul style="list-style-type: none">~Minimum of two alphabetic characters<ul style="list-style-type: none">>one uppercase>one lowercase~minimum of one numeric/special character(s)	<ul style="list-style-type: none">~No more than two repetitive/sequential characters (for example: ll, aa, 12, or 45 are acceptable while lll, aaa, 123, or 456 are not acceptable)~Do not use dictionary words, usernames, or context specific words (for example, INL, Idaho, Nuclear, HPC, etc.)>A phrase composed of dictionary words is acceptable	<ul style="list-style-type: none">~Do not reuse any of the previous 24 passwords associated with the account.~Never reuse compromised passwords	<ul style="list-style-type: none">~ Passwords should be different from passwords used for other accounts~ A passphrase is easier to remember~ Utilize a password manager to remember your password~do not share your password with anyone (HPC Staff will never ask for you password)

HPC Geographical Access Policy

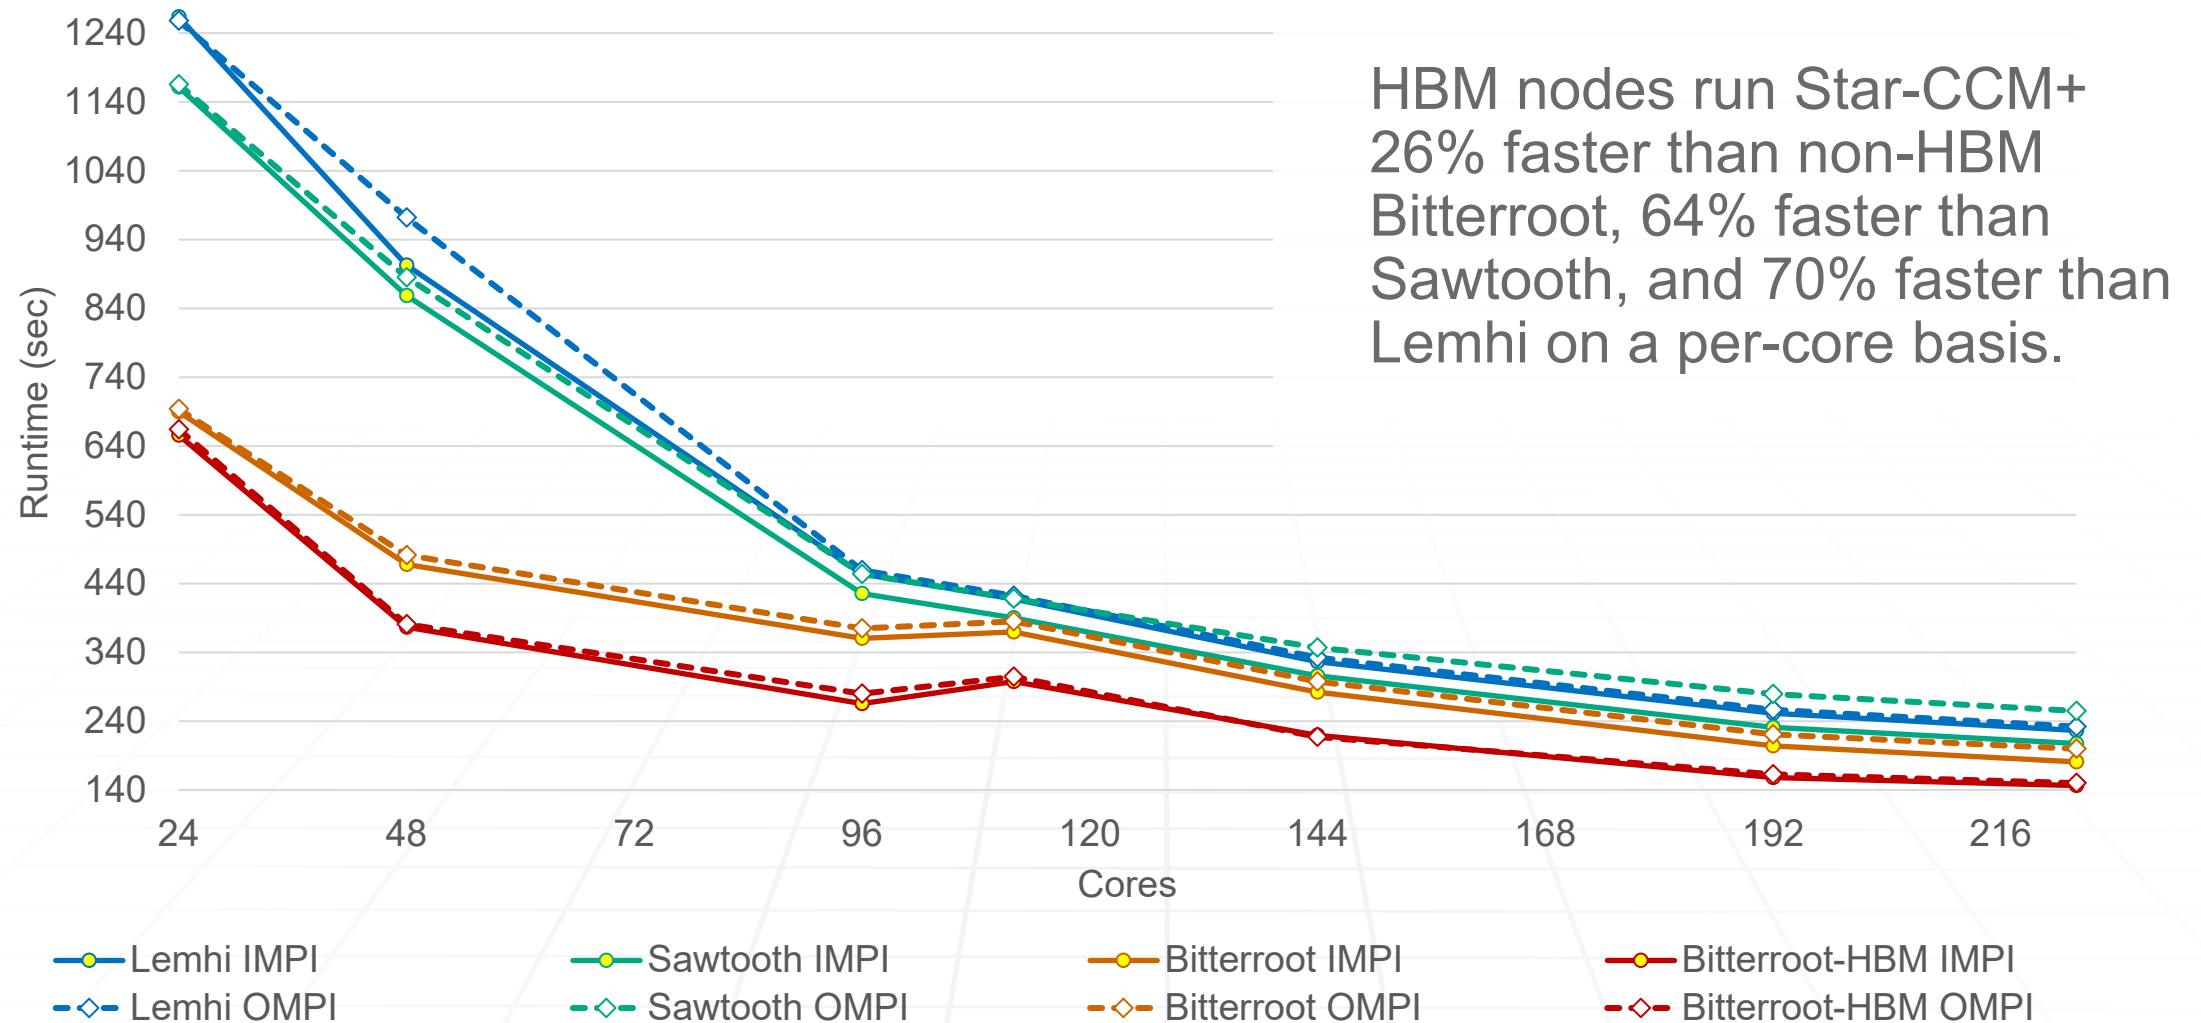
- **Purpose:** to establish clear guidelines regarding the geographical locations from which users are authorized to access HPC system resources.
- This policy applies to all users of HPC resources (both internal and external)



Fairshare & Limits Policy

- Both PBS and Slurm are configured to balance fair-share (a measure of recent resource usage) and job age (amount of time a job has been waiting to run).
 - Fair-share is based off a half-life formula that favors the most recent usage statistics.
 - In PBS, all jobs accrue “age priority” once they are submitted to the queue.
 - In Slurm, we have it configured such that only the oldest 10 queued jobs for each user accrue age priority.
- We also have some queue/partition limits in place to prevent individual users from monopolizing the system.
 - Sawtooth general queue: up to 1368 nodes per user (queued and running)
 - Lemhi general queue: up to 384 nodes per user (queued and running)
 - Bitterroot general partition: up to 168 nodes per user (running)
 - See <https://hpcweb.hpc.inl.gov/tools/ClusterJobInfo>

Star- CCM Results



C3 Data Center Outage

- **Date of outage:** October 7-11, 2024
- 7 Oct Full outage for datacenter to support new UPS tie-in
 - No people in building
 - HPC systems off
- 8-9 Oct medium outage
 - People can work in the building but HPC systems off
 - Row 9,10 working. No UPS
- 10-11 Oct – continued medium outage
 - UPS vendor will commission and validate UPS system
 - Once complete, all HPC systems can be powered on

Script Generator 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
```

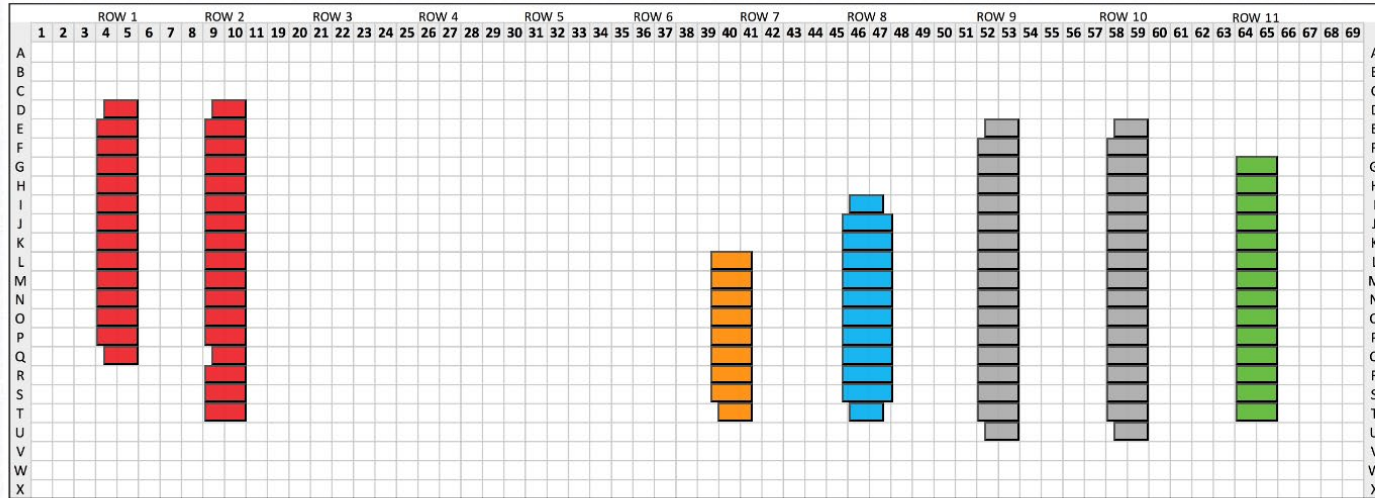
Lemhi Decommission Date

- Lemhi will be turned over to the Idaho State Board of Education on April 1, 2025
- After that date
 - Lemhi will not be approved for INL data
 - Lemhi will not be managed by INL HPC
 - No further software support will be provided for Lemhi
- The Wind River system (94,416 cores) will be operational before the turn-over
- Teton (400k cores) will be operational on or before 1 July 2025

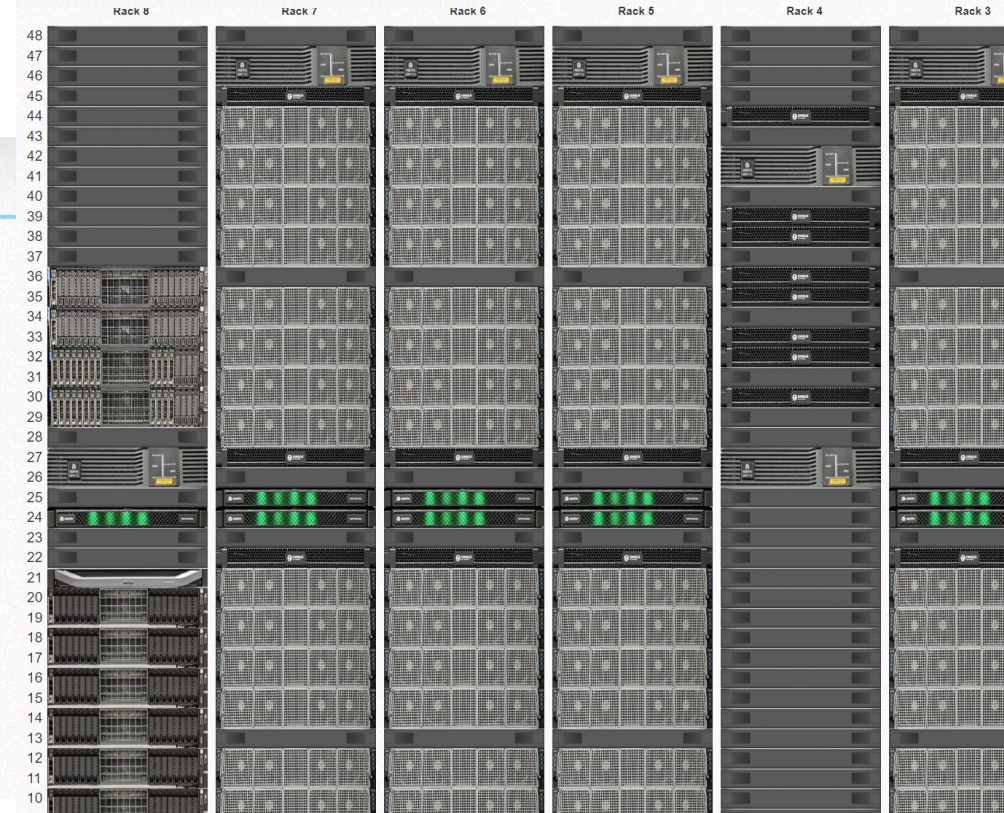
HPC Data Center Visual Introduction

- Interactive view of the C3 Datacenter

DATA CENTER VISUAL



BITTERROOT PHYSICAL FRONT



Physical Status Page

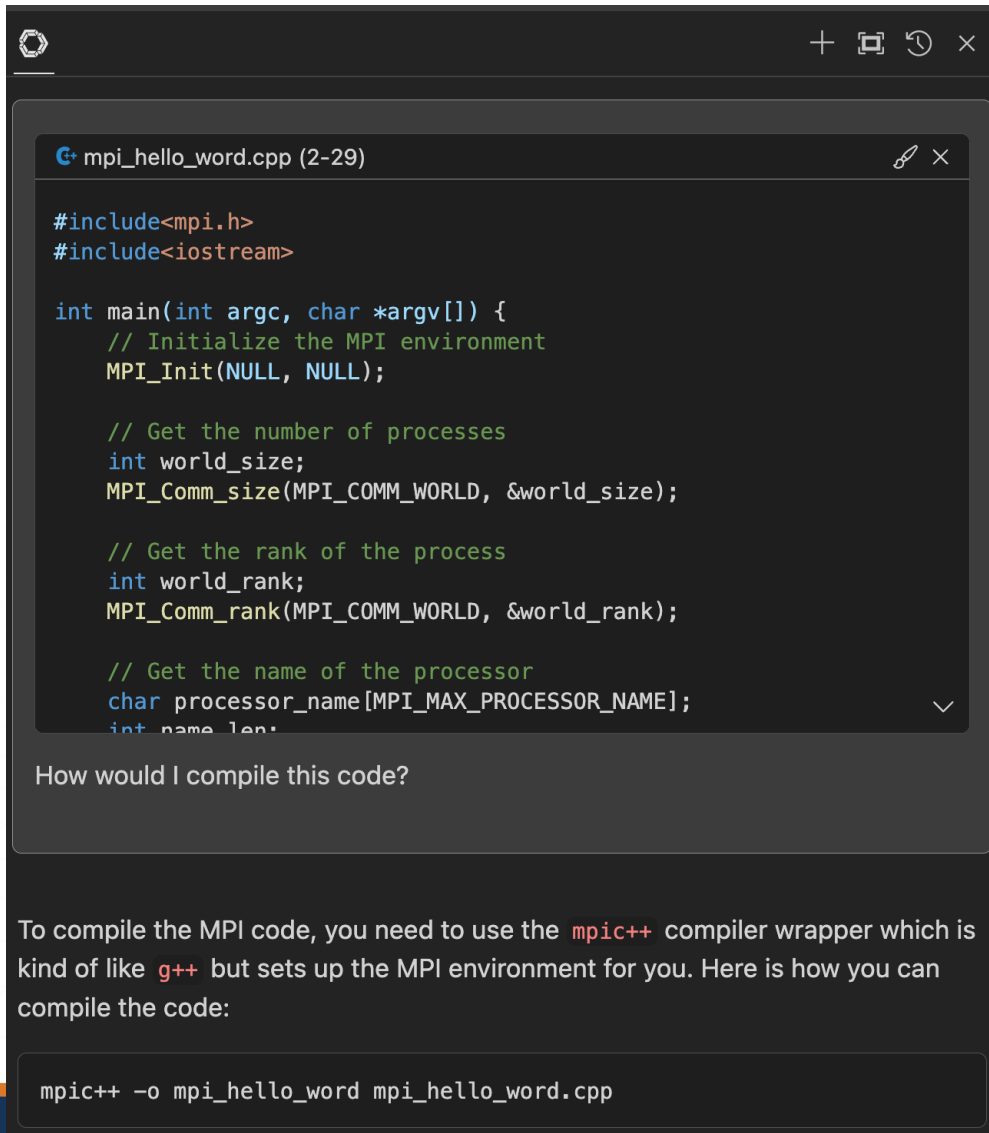
- Default view is the current CPU load percent
- Reached through the job status page
- Can search by username to see what nodes a user is running on

Type	Job Count	Description
Job Exclusive	266	Job is running with all cores or job has taken exclusive use of node.
Free	113	Node has 0 cores in use.
Offline	0	Node is responsive but is not available for use.
Unknown State	0	Node state is undeterminable by PBS.
Down	5	Node is completely powered down or is unlocatable.

Type	Job Count	Description
Normal Load	336	Node has an acceptable load average for the jobs running.
Abnormal Load on Empty Node	42	Node has an abnormally high load for zero jobs.
Low Load Node	6	Node has an abnormally low load for the cores requested.
High Load Node	0	Node has an abnormally high load for the cores requested.

Rack 4	
38	25
12	38
38	0
25	0
38	13
12	25
38	25
25	50
43	44
43	41
42	57
44	44
0	43
12	46
45	45
43	46
88	89
89	88
86	88
87	88
87	86
90	86
86	86
88	89
25	25
12	38
25	12
25	25
0	50
0	25
13	12
38	38

VSCode Co-pilot



The screenshot shows the VS Code editor with a file named `mpi_hello_word.cpp` (lines 2-29). The code is a simple MPI "Hello, World!" program. Below the code editor, a chat window is open with the prompt "How would I compile this code?". The response explains that the `mpic++` compiler wrapper should be used and provides the command `mpic++ -o mpi_hello_word mpi_hello_word.cpp`.

```
#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);

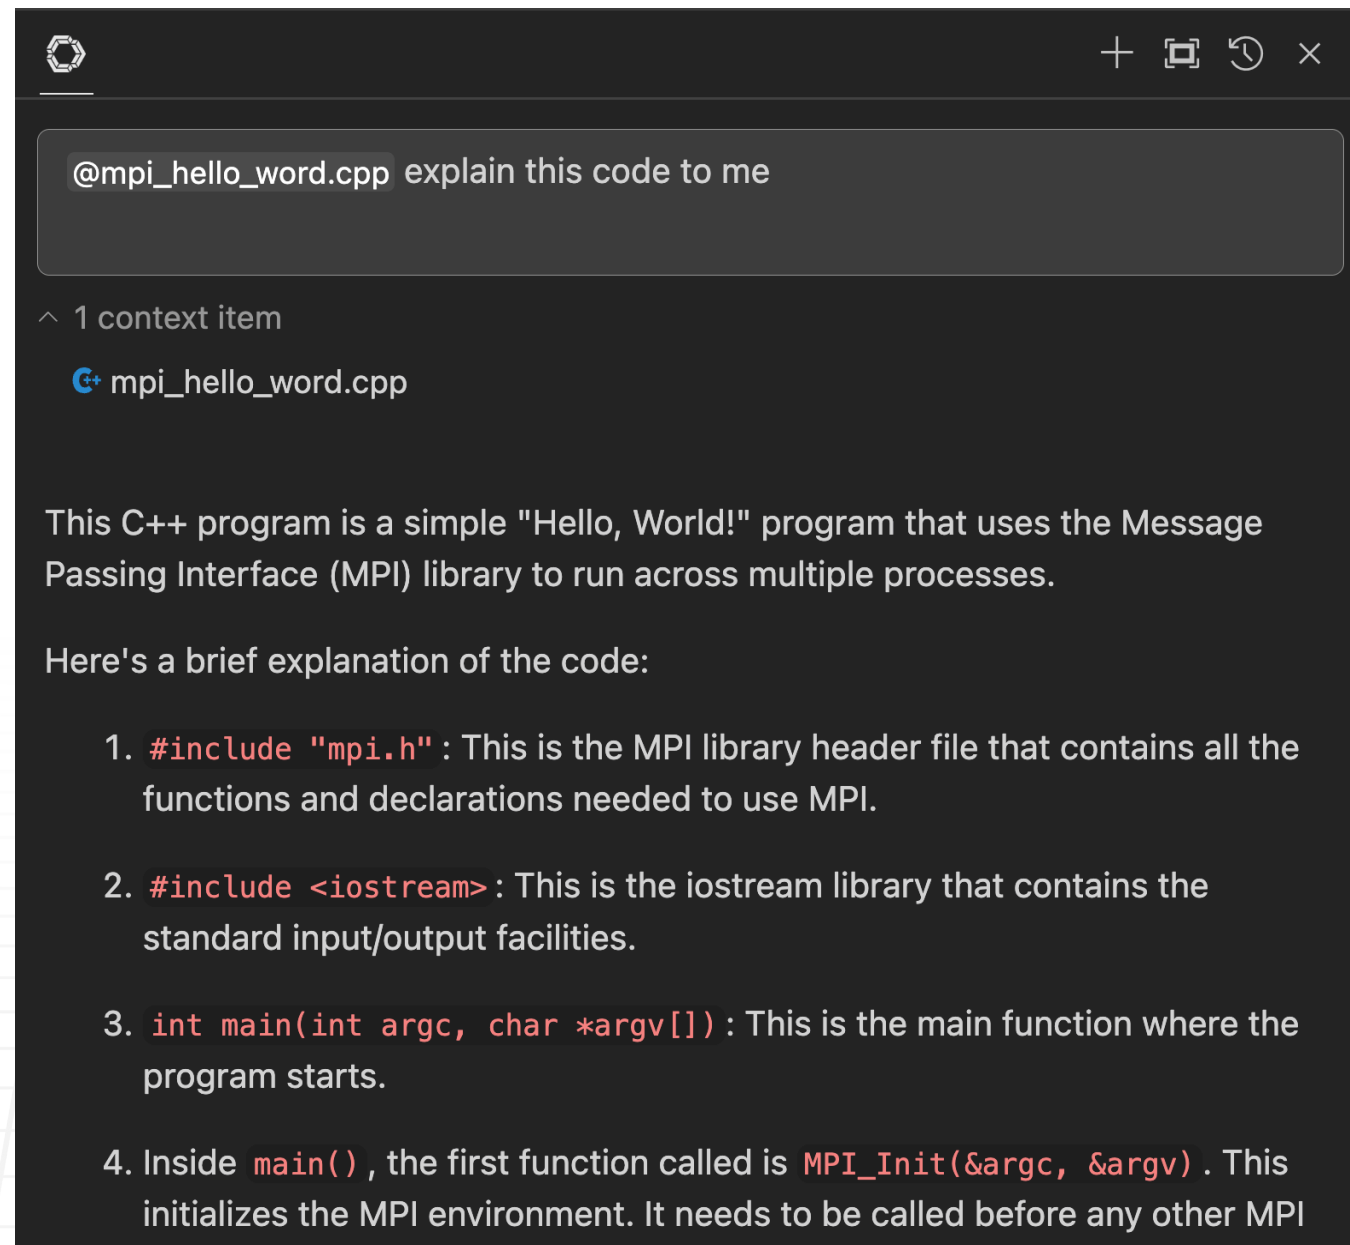
    // Get the rank of the process
    int world_rank;
    MPI_Comm_rank(MPI_COMM_WORLD, &world_rank);

    // Get the name of the processor
    char processor_name[MPI_MAX_PROCESSOR_NAME];
    int name_len;
```

How would I compile this code?

To compile the MPI code, you need to use the `mpic++` compiler wrapper which is kind of like `g++` but sets up the MPI environment for you. Here is how you can compile the code:

```
mpic++ -o mpi_hello_word mpi_hello_word.cpp
```



The screenshot shows the VS Code editor with a file named `mpi_hello_word.cpp`. The chat window shows the prompt "@mpi_hello_word.cpp explain this code to me". The response includes a context item for the file and a detailed explanation of the code's purpose and components.

@mpi_hello_word.cpp explain this code to me

^ 1 context item

`mpi_hello_word.cpp`

This C++ program is a simple "Hello, World!" program that uses the Message Passing Interface (MPI) library to run across multiple processes.

Here's a brief explanation of the code:

1. `#include "mpi.h"`: This is the MPI library header file that contains all the functions and declarations needed to use MPI.
2. `#include <iostream>`: This is the iostream library that contains the standard input/output facilities.
3. `int main(int argc, char *argv[])`: This is the main function where the program starts.
4. Inside `main()`, the first function called is `MPI_Init(&argc, &argv)`. This initializes the MPI environment. It needs to be called before any other MPI

Training Cluster Update

- New hardware has been purchased replacing 11-year-old hardware
- 15 nodes
 - 10 CPU/GPU nodes
 - 5 CPU only nodes
 - 2,880 AMD "Genoa" cores
 - 20 NVIDIA L4 GPUs
- Will run Open OnDemand
- Availability for training is still in progress.

INL Open OnDemand (HPC OnDemand) Desktop vs. Desktop with GPU

Interactive Apps

Desktops

Linux Desktop

Linux Desktop with GPU

IDE

VSCode Desktop

VSCode Server

Jupyter

Jupyter

RDM

Research Data Management

NCRC

GUIs

NEAMS Workbench

Herd

Code Execution

Tests

Build Test Suite

Training Videos

Linux Desktop version: 4ee3ad4

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](#).

Interactive Apps

Desktops

Linux Desktop

Linux Desktop with GPU

IDE

VSCode Desktop

VSCode Server

Jupyter

Jupyter

RDM

Research Data Management

NCRC

GUIs

NEAMS Workbench

Herd

Code Execution

Tests

Build Test Suite

Training Videos

Bison Videos

Relap 5

Training

Training Sessions

NEAMS Griffin Neutronics Fast Reactor Modeling and Simulation

Shift/Titan Training

Linux Desktop with GPU version: 4ee3ad4

This app will launch an interactive desktop on a visualization node.

This is meant for all types of tasks such as:

- accessing & viewing files
- compiling code
- debugging
- running visualization software **with** 3D hardware support

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

Viz

Number of Hours

1

Warning max walltime could be different between systems. Please see queues section on the [cluster queues](#) for more information.

Cluster	Max Hours
Sawtooth	12
Viz	72

GPUs Requested

1

Min 1 | Max 2

Launch

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

INL Open OnDemand (OnDemand) Linux Desktop

Interactive Apps

AI/ML Training

AI/Machine Learning Tutorial Series

MIT Symposium Summer 2021

Active Trainings

Shift/Titan Training

Desktops

Linux Desktop

GUIs

Barracuda VR

IDE

VSCode Container

VSCode Desktop

VSCode Server

Jupyter

Jupyter

NEAMS Workbench

Linux Desktop

This app will launch a Linux Desktop in the INL HPC enclave.

Project

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

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 without hardware acceleration
Hardware Accelerated	Running visualization software with 3D hardware acceleration

Cluster

Lemhi

Select what cluster you want to run your desktop on.

Launch

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

Interactive Apps

AI/ML Training

AI/Machine Learning Tutorial Series

MIT Symposium Summer 2021

Active Trainings

Shift/Titan Training

Desktops

Linux Desktop

GUIs

Barracuda VR

IDE

VSCode Container

VSCode Desktop

VSCode Server

Jupyter

Jupyter

NEAMS Workbench

Linux Desktop

This app will launch a Linux Desktop in the INL HPC enclave.

Project

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

Job Type

Hardware Accelerated

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 without hardware acceleration
Hardware Accelerated	Running visualization software with 3D hardware acceleration

Cluster

Viz

Select what cluster you want to run your desktop on.

Number of Hours

72

Warning, max walltime could be different between systems. Please see queues section on the [cluster queues](#) for more information.

Cluster	Max Hours*
Bitterroot	72
Lemhi	72
Sawtooth	168
Viz	72

* The number of hours assigned to your session may be reduced if a planned outage is forthcoming.

GPUs Requested

0

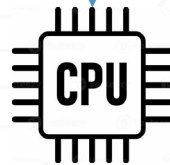
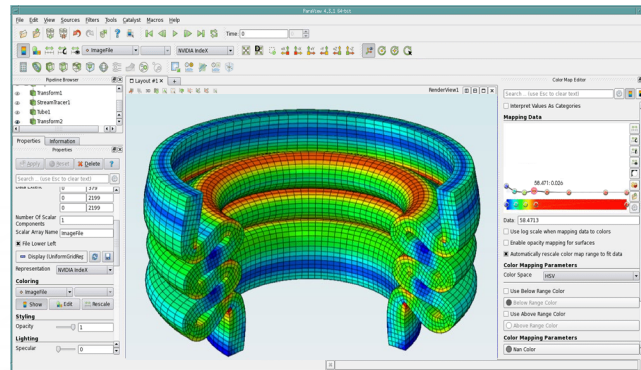
Min 0 | Max 2. Requesting GPUs automatically updates the amount of CPUs requested.

Launch

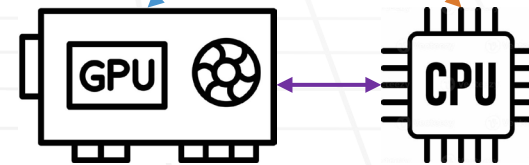
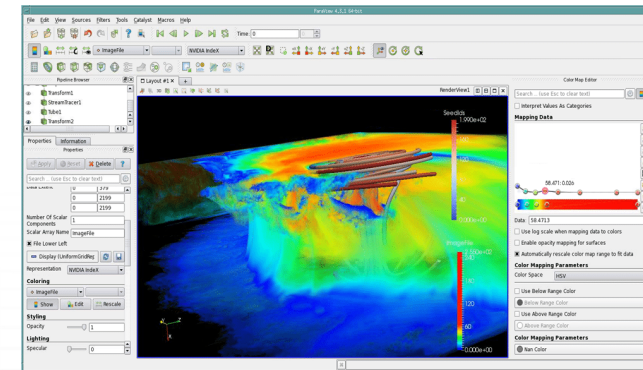
* The Linux Desktop session data for this session can be accessed under the [data root directory](#).

INL Open OnDemand (HPC OnDemand and OnDemand) Desktop vs. Desktop with GPU or Linux Desktop

Desktop



Desktop with GPU



Arbiter

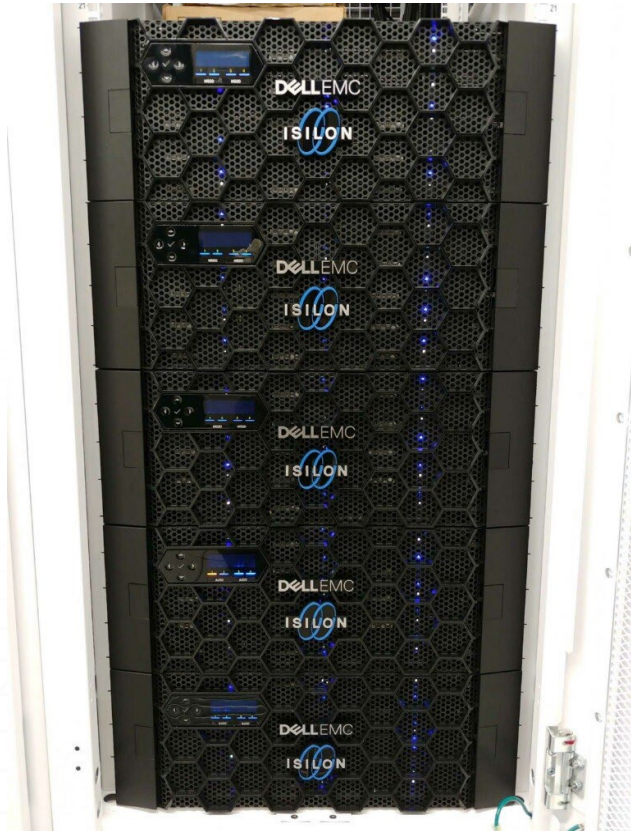


- Arbiter2
 - University of Utah Center for High Performance Computing
 - [GitHub - chpc-uofu/arbiter2: A daemon that uses cgroups to monitor and manage user behavior on login nodes](https://github.com/chpc-uofu/arbiter2)
- Monitors and protects interactive nodes with cgroups (login nodes)
 - Records the activity on all nodes where installed and running
 - Automatically sets limits on the resources available to each user
 - Notifies users by email when excessive resources use is detected
- Three levels of notifications, resource reduction
 - Level 1 – 80% of starting resources for a time period (30 minutes)
 - Level 2 – 50% of starting resources for a time period (1 hour)
 - Level 3 – 30% of starting resources for a time period (2 hours)

Wind River Update

- Schedule
 - Sept 13 – Rack build complete at Dell (sub)
 - Sept 23 – Scheduled arrival of system
 - TBD – rear door heat exchangers, on chip cooling manifolds, dell cabling
 - 2-4 Weeks – HPC testing and configuration
 - Late November – aspirational availability
- 20 racks – 17 compute
- room for 1120 compute nodes – 843 nodes initially (94,416 cores)
- Nodes are the same as Bitterroot
 - 112 cores (2x Intel(R) Xeon(R) Platinum 8480+)
 - 256GB memory
 - 200Gb/s Omni-Path fabric

Isilon Home and Project space



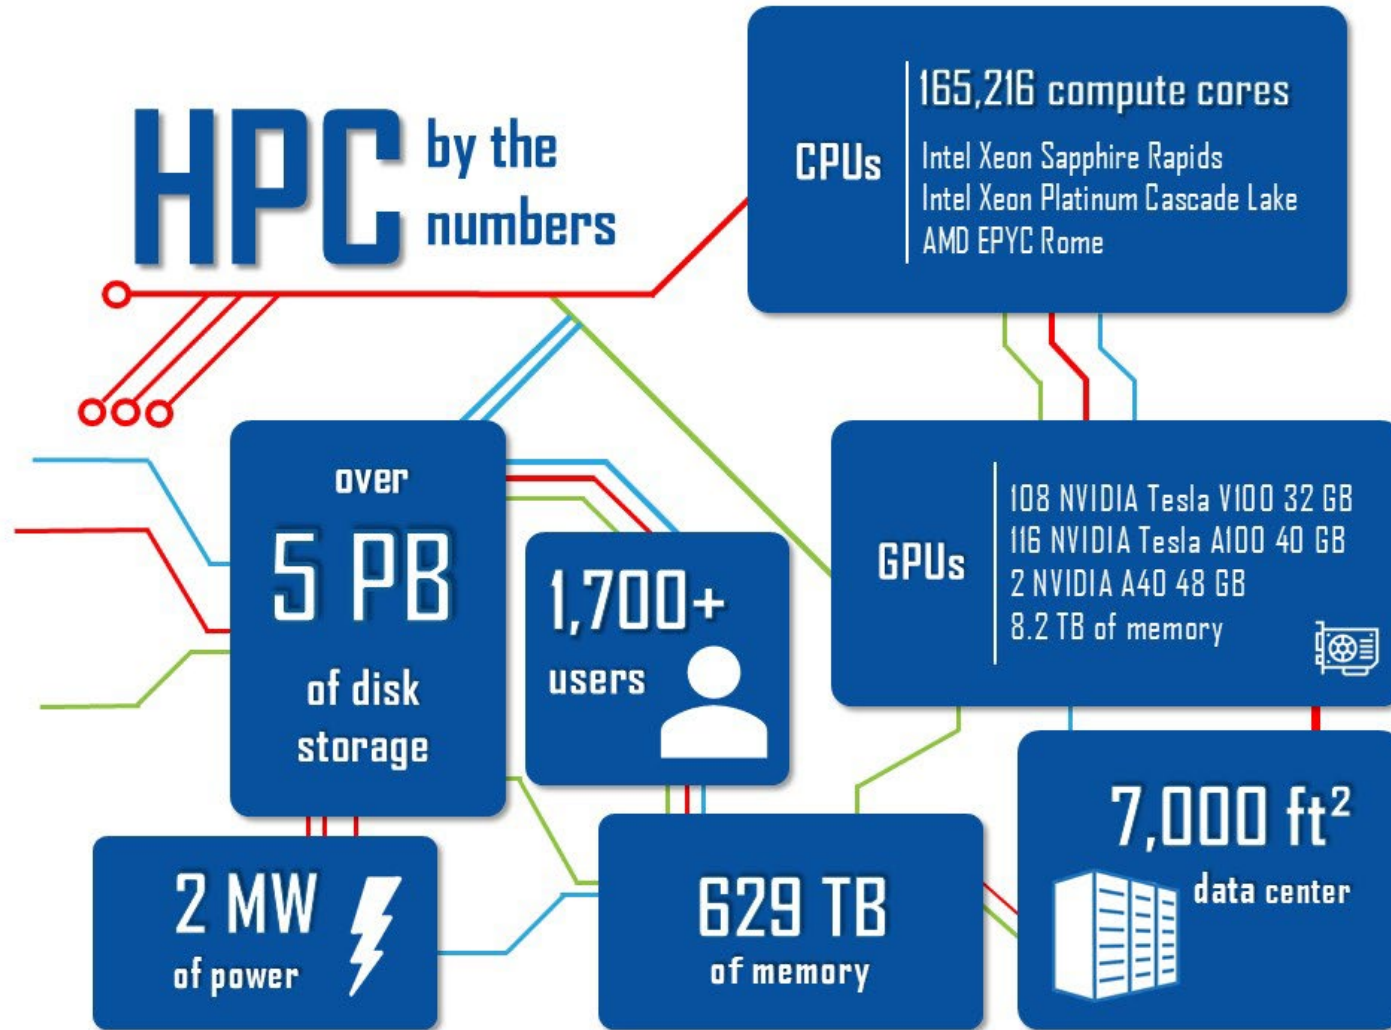
- Was able to get down to 74% full - from 82%
- The project space is taking **82%**
- Newly created projects will have a 20TB quota

IBM ESS Scratch



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

HPC by the Numbers Graphic Update



Next HPC User Group Meeting

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

Questions

- Please type your questions into chat