

22 May 2024

Matthew Anderson
Manager High Performance Computing

NRDS and NSUF Data Management

Nuclear Science User Facilities

High Performance Computing Resources

- NSUF HPC systems support a wide range of users and programs as a shared-use resource for national laboratories, universities, and industry
- Sawtooth (2020)
  - 6 Petaflops performance
  - 2,079 compute nodes, 99,972 compute cores
  - #37 on November 2019 TOP500 list
- Lemhi (2018)
  - 1 Petaflop performance
  - 504 compute nodes, 20,160 compute cores
  - #427 on November 2018 TOP500 list

Sawtooth



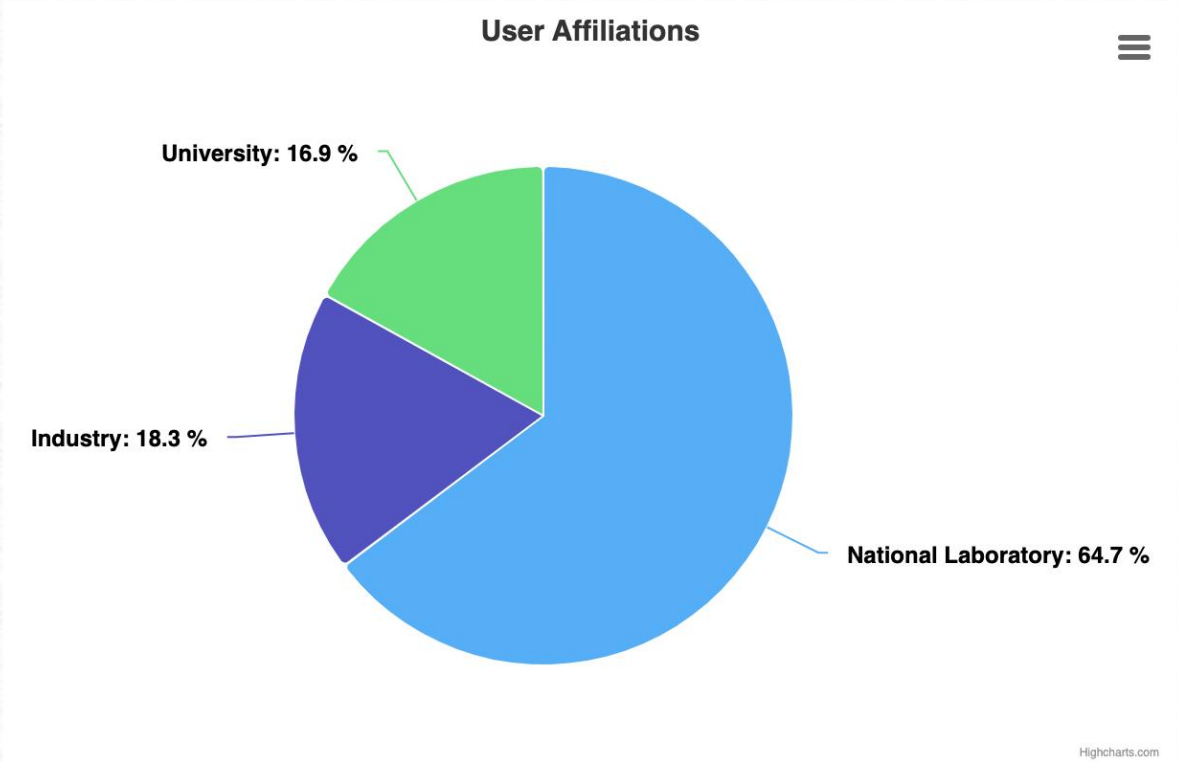
Lemhi

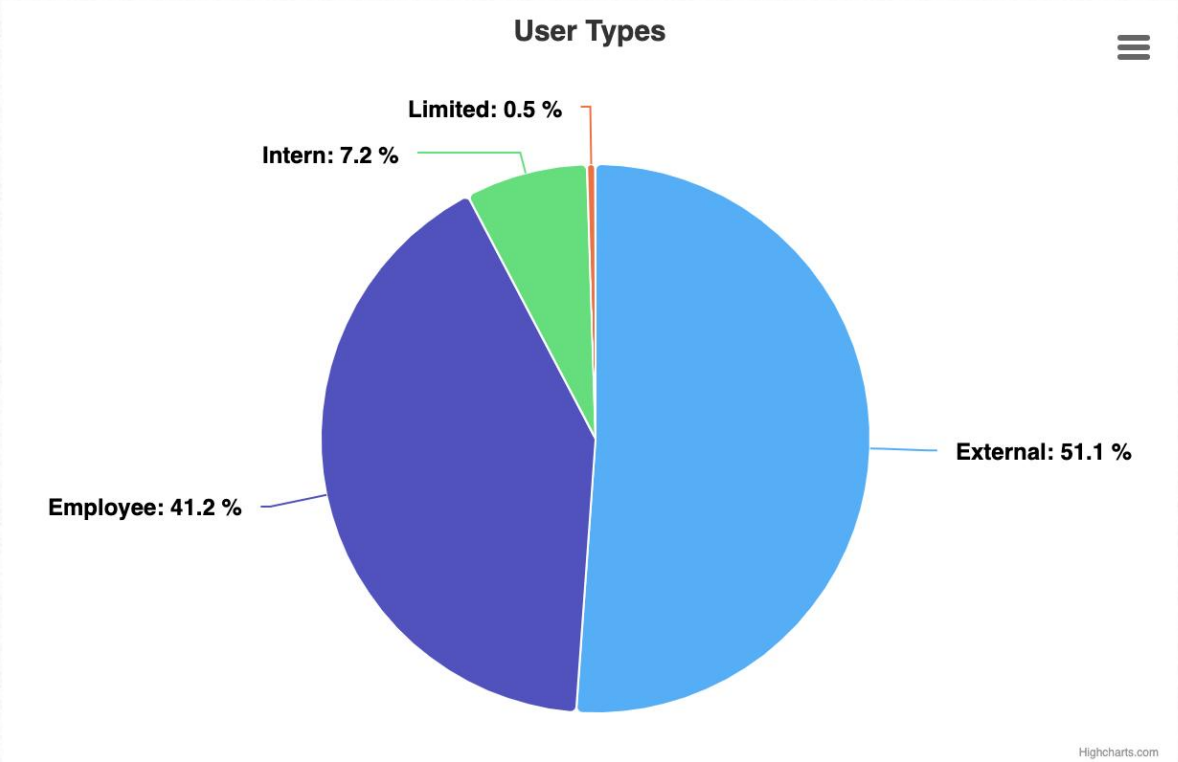


A right-sized solution for DOE Nuclear Energy research and development

HPC users on 31 July 2023

Total Users: 1141 - View Details

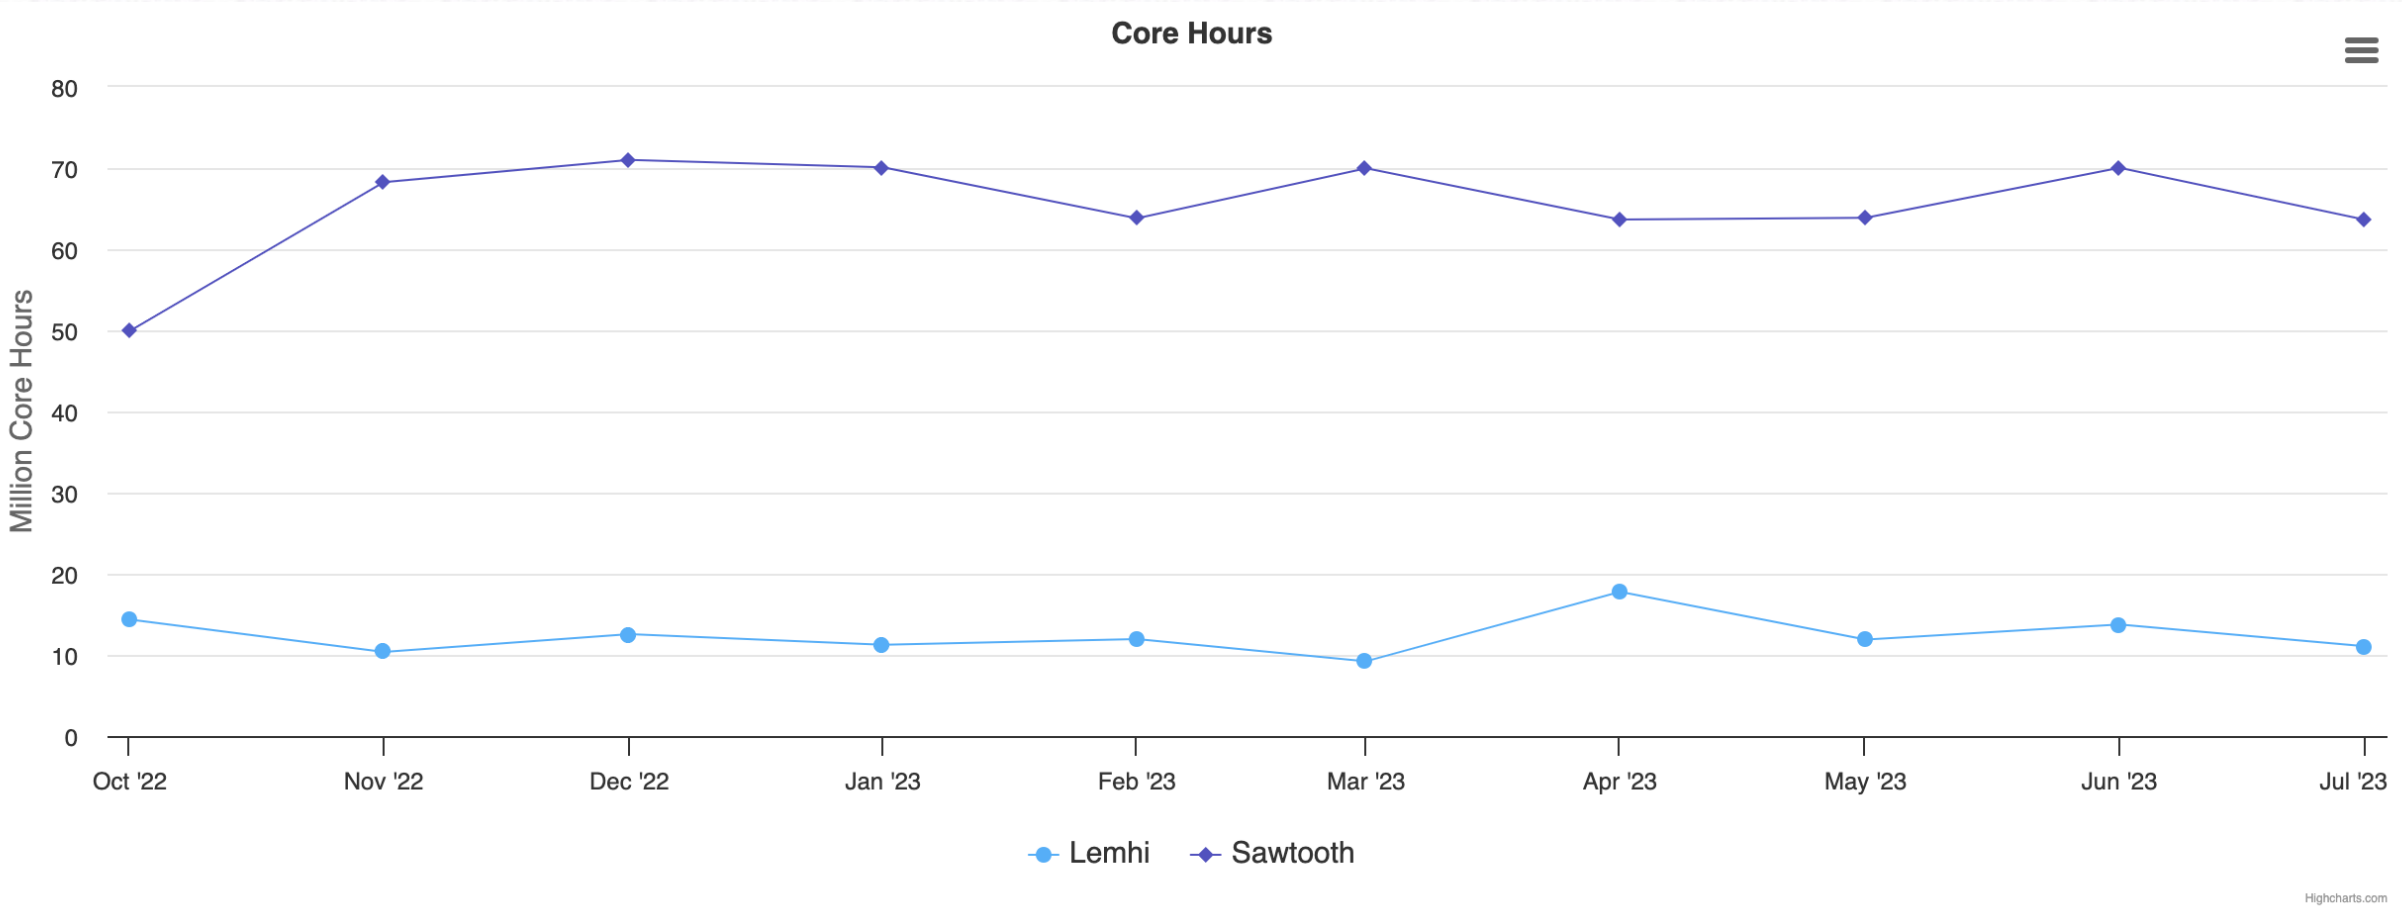




Currently supporting users from 8 National Laboratories, 41 universities and 35 companies.

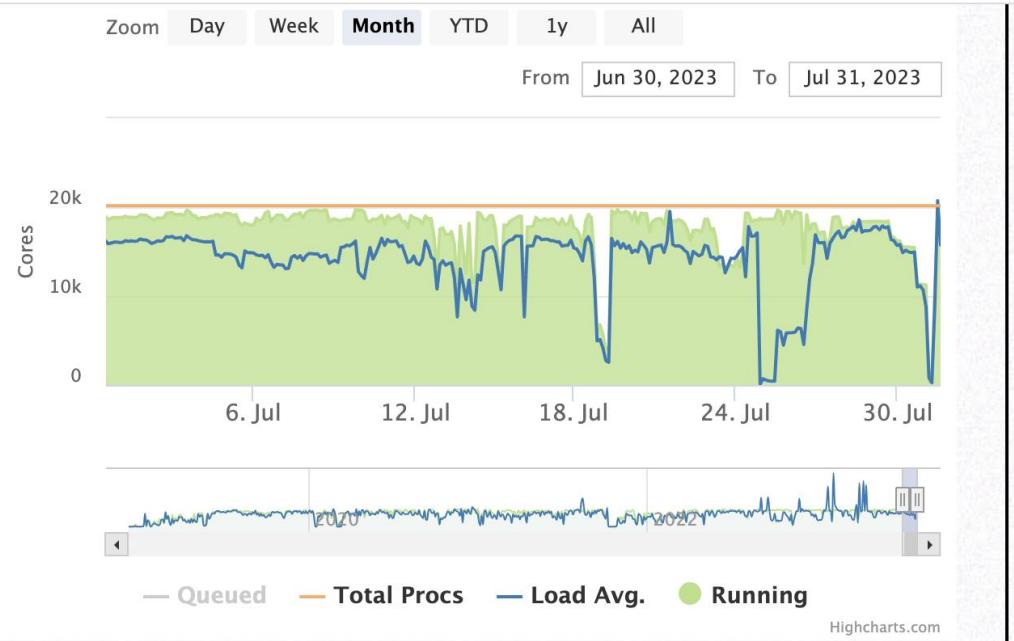
International institutions supported are from Belgium, Canada, Czech Republic, Denmark, Finland, France, Indonesia, Italy, Japan, Jordan, Poland, Slovakia, Spain, and the United Kingdom.

Sawtooth and Lemhi delivered core-hours per month over the fiscal year





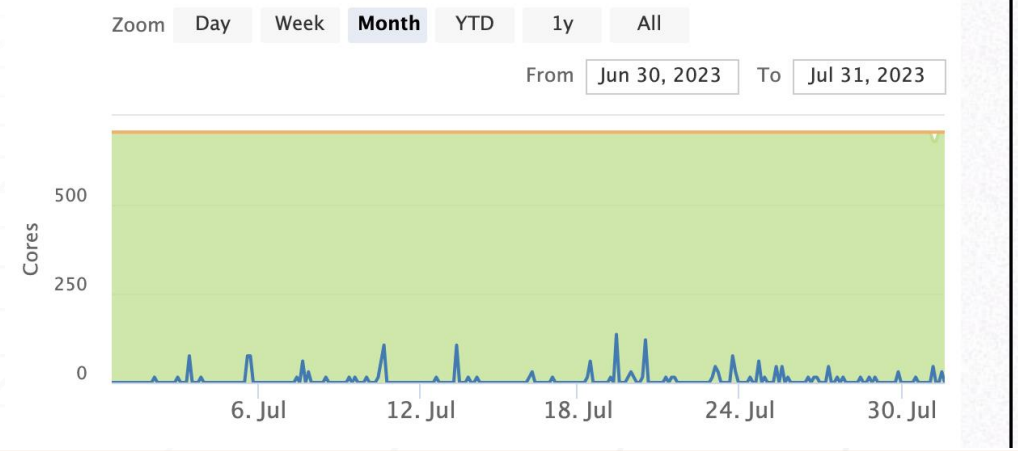
Max Running: 99,792 | Max Load: 179,012 | Max Queued: 1,344,877
Avg Running: 97,443 | Avg Load: 70,521 | Avg Queued: 805,377



Max Running: 19,839 | Max Load: 23,285 | Max Queued: 3,704,840
Avg Running: 17,612 | Avg Load: 14,139 | Avg Queued: 1,644,846

oodoo

Last Updated: 7/31/2023, 1:30:31 PM

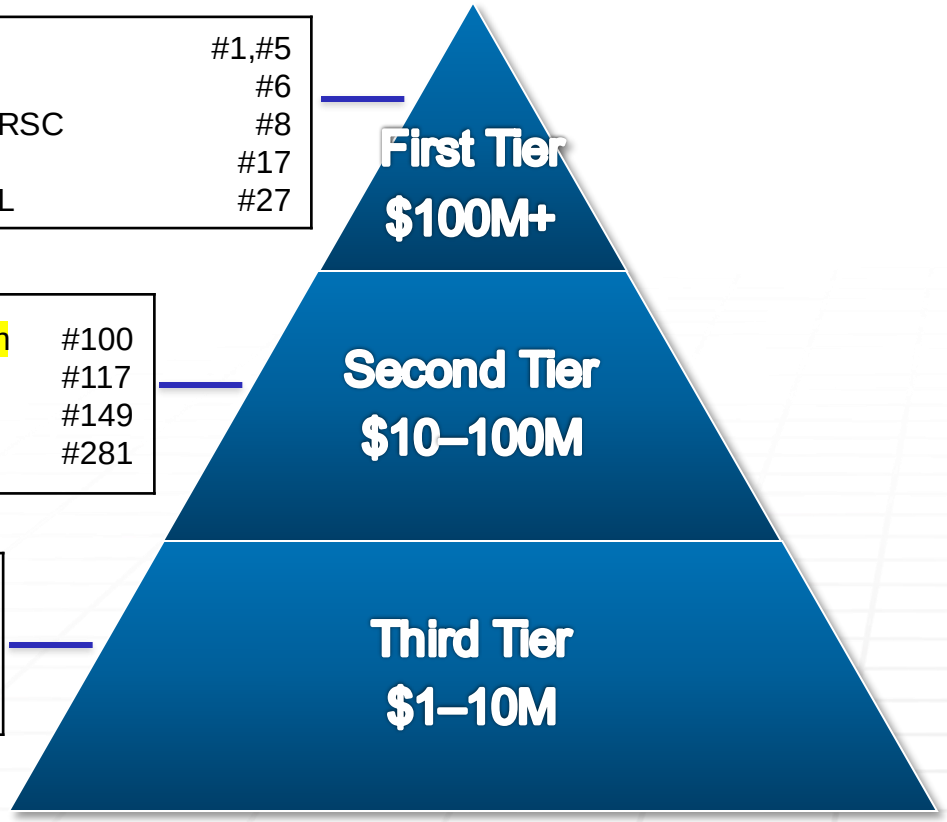


DOE Ecosystem – HPC Trends – INL positioning

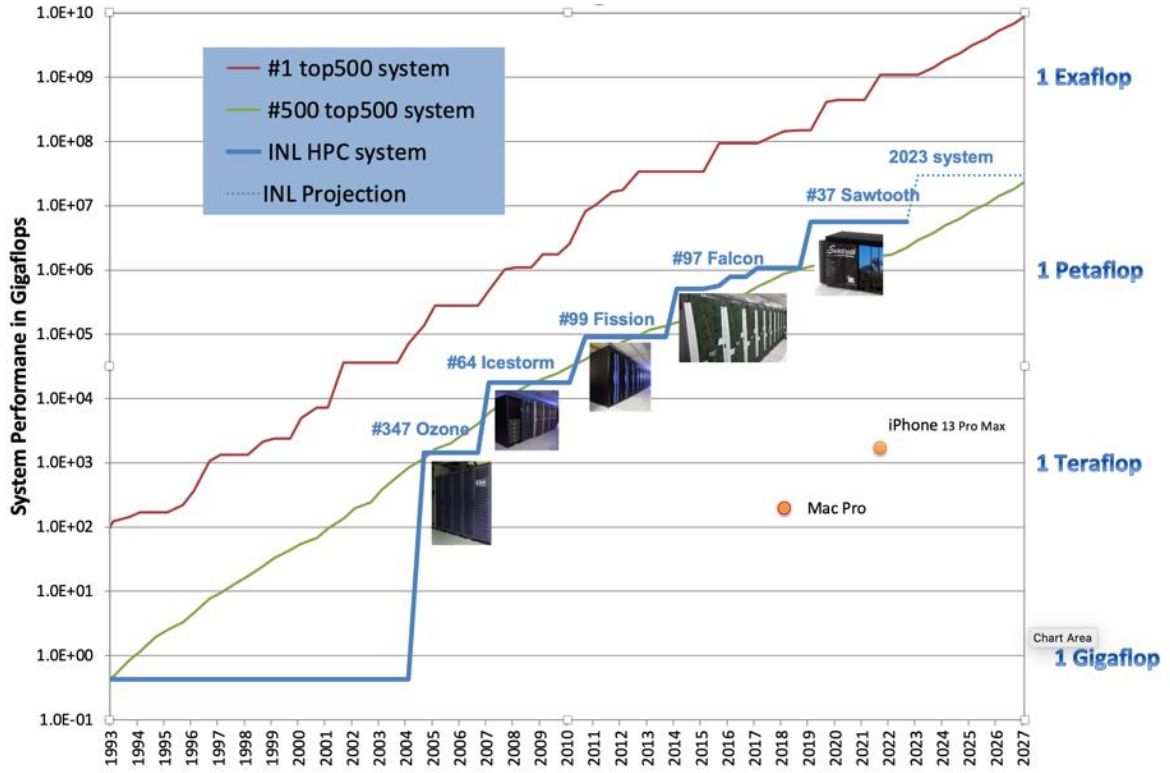
ORNL	#1,#5
LLNL	#6
LBNL/NERSC	#8
ANL	#17
LANL/SNL	#27

INL Sawtooth	#100
NREL	#117
NETL	#149
PNNL	#281

INL Lemhi	#500+
Bettis	#500+
Knolls	#500+

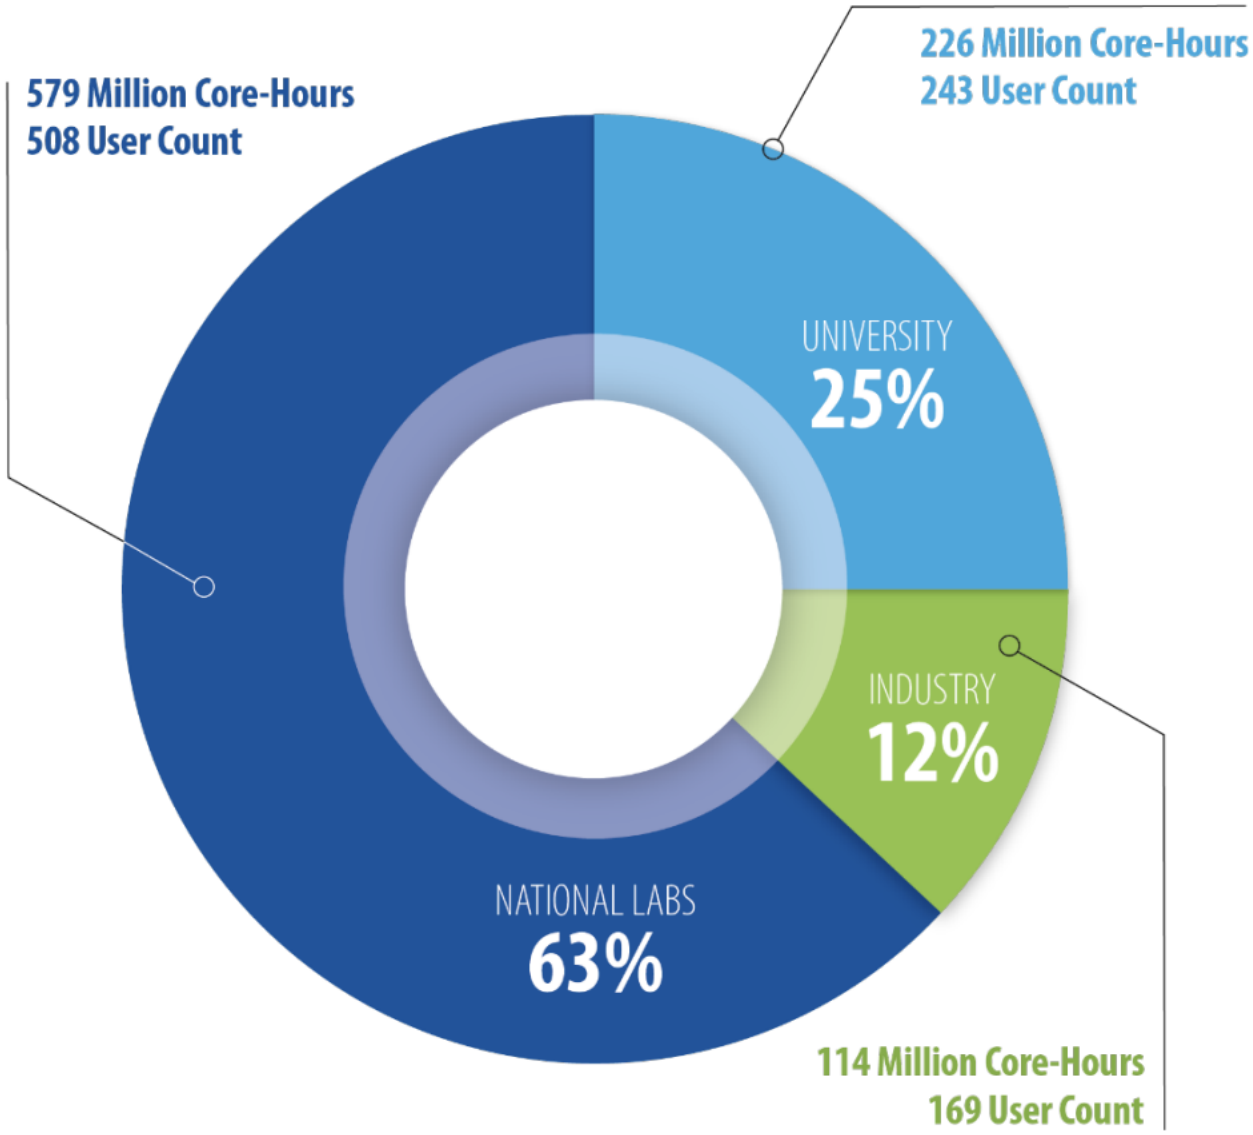


November 2022 Top500 Rankings



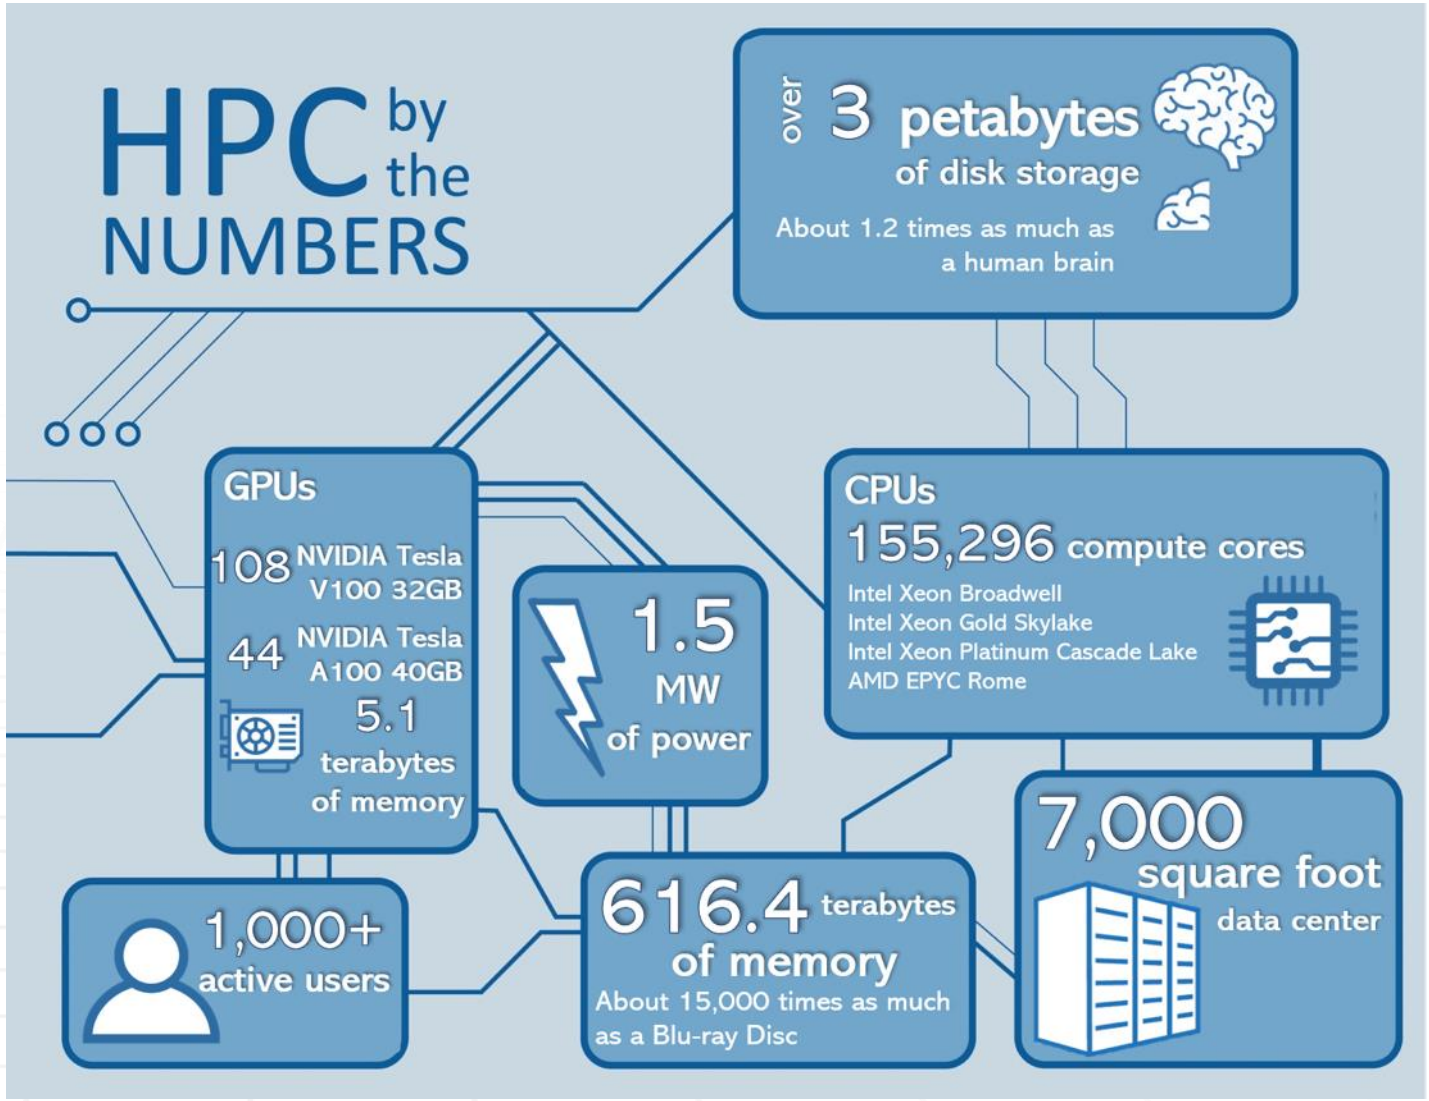
Continual updates are required: ~20–30PF for 2023 system

HPC Utilization for FY-22



HPC Systems

- Sawtooth
- Lemhi (CPU)
- Hoodoo (GPU)
- Viz
- ARM
- Storage
  - EMC Isilon
  - IBM GPFS



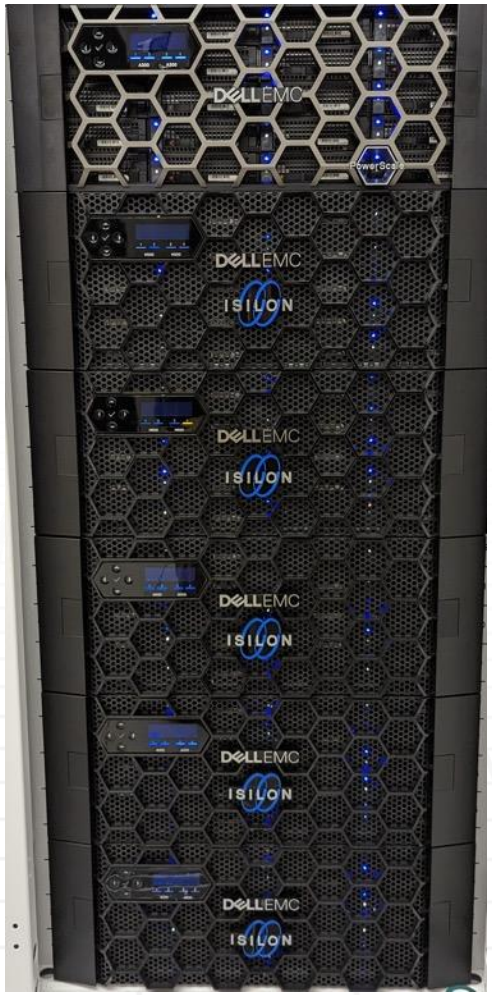
HPC Systems – Isilon/GPFS/Pure



- IBM GPFS
  - 3 PB
  - NVME/HDD
  - /scratch



- Pure
  - All-flash
  - /apps



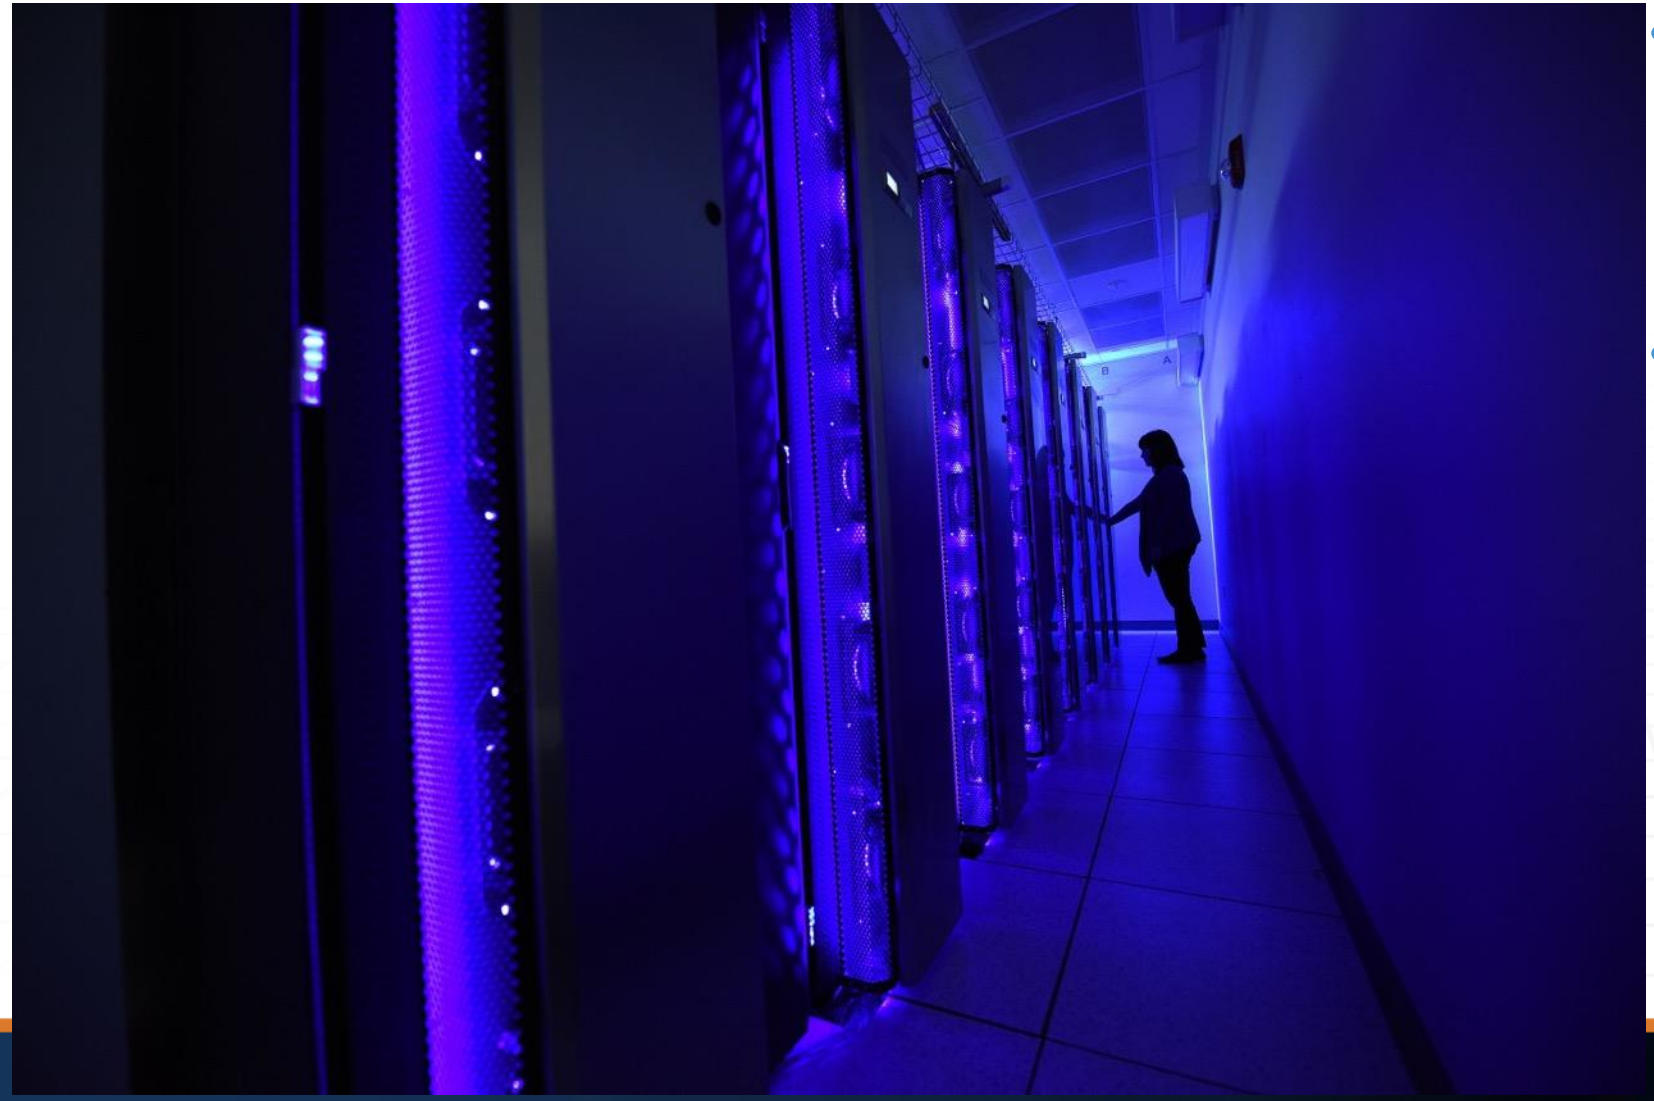
- EMC Isilon
  - 2.1 PB
  - HDD
  - /home
  - /projects
- NRDS
  - 1.7 PB smartlock

HPC Systems – Sawtooth



- Nodes
  - 4 Login
  - 8 Gateway
  - 2,079 Compute
- Specs
  - 99,792 Intel Cascade Lake cores
  - 395 TB Total Memory
  - 108 V100 32GB GPUs
  - EDR 100 Gb/s InfiniBand
    - Enhanced Hypercube

HPC Systems – Lemhi



- Nodes
  - 2 Login
  - 4 Gateway
  - 504 Compute
- Specs
  - 20,160 Intel Sky Lake cores
  - 94.5 TB total memory
  - Cornelis 100 Gb/s OmniPath
    - Fat-tree

HPC Systems – Hoodoo



- Nodes
  - 1 Login
  - 11 Compute
- Specs
  - 352 AMD EPYC Rome cores
  - 5.5 TB Total Memory
  - 44 A100 40GB GPUs
  - HDR 200 Gb/s InfiniBand
    - Fat-tree

HPC Systems – Viz



- Nodes
  - 1 Login
  - 5 Compute
- Specs
  - 88 Intel Xeon cores
  - 1.4 TB Total Memory
  - GPUs
    - 8 Quadro M6000 24GB
    - 4 Quadro K6000 12GB
    - 2 Quadro P6000 24GB
    - 2 Quadro P5000 16GB
  - 40 Gb/s Ethernet
    - Fat-tree

HPC Systems – ARM

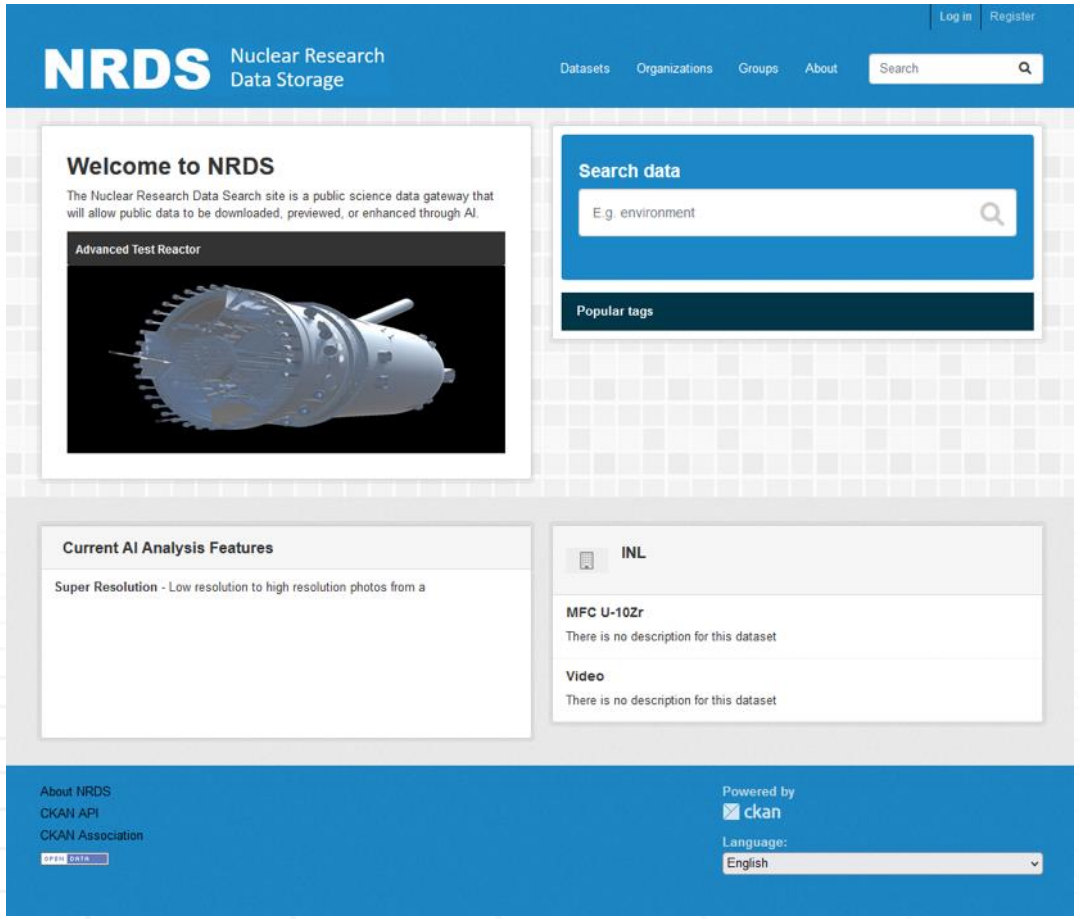


- Nodes
  - 1 Login
  - 16 Compute
- Specs
  - 768 A64FX cores
  - 512 GB High Bandwidth Total Memory
  - EDR 100Gb/s InfiniBand
    - Fat-tree

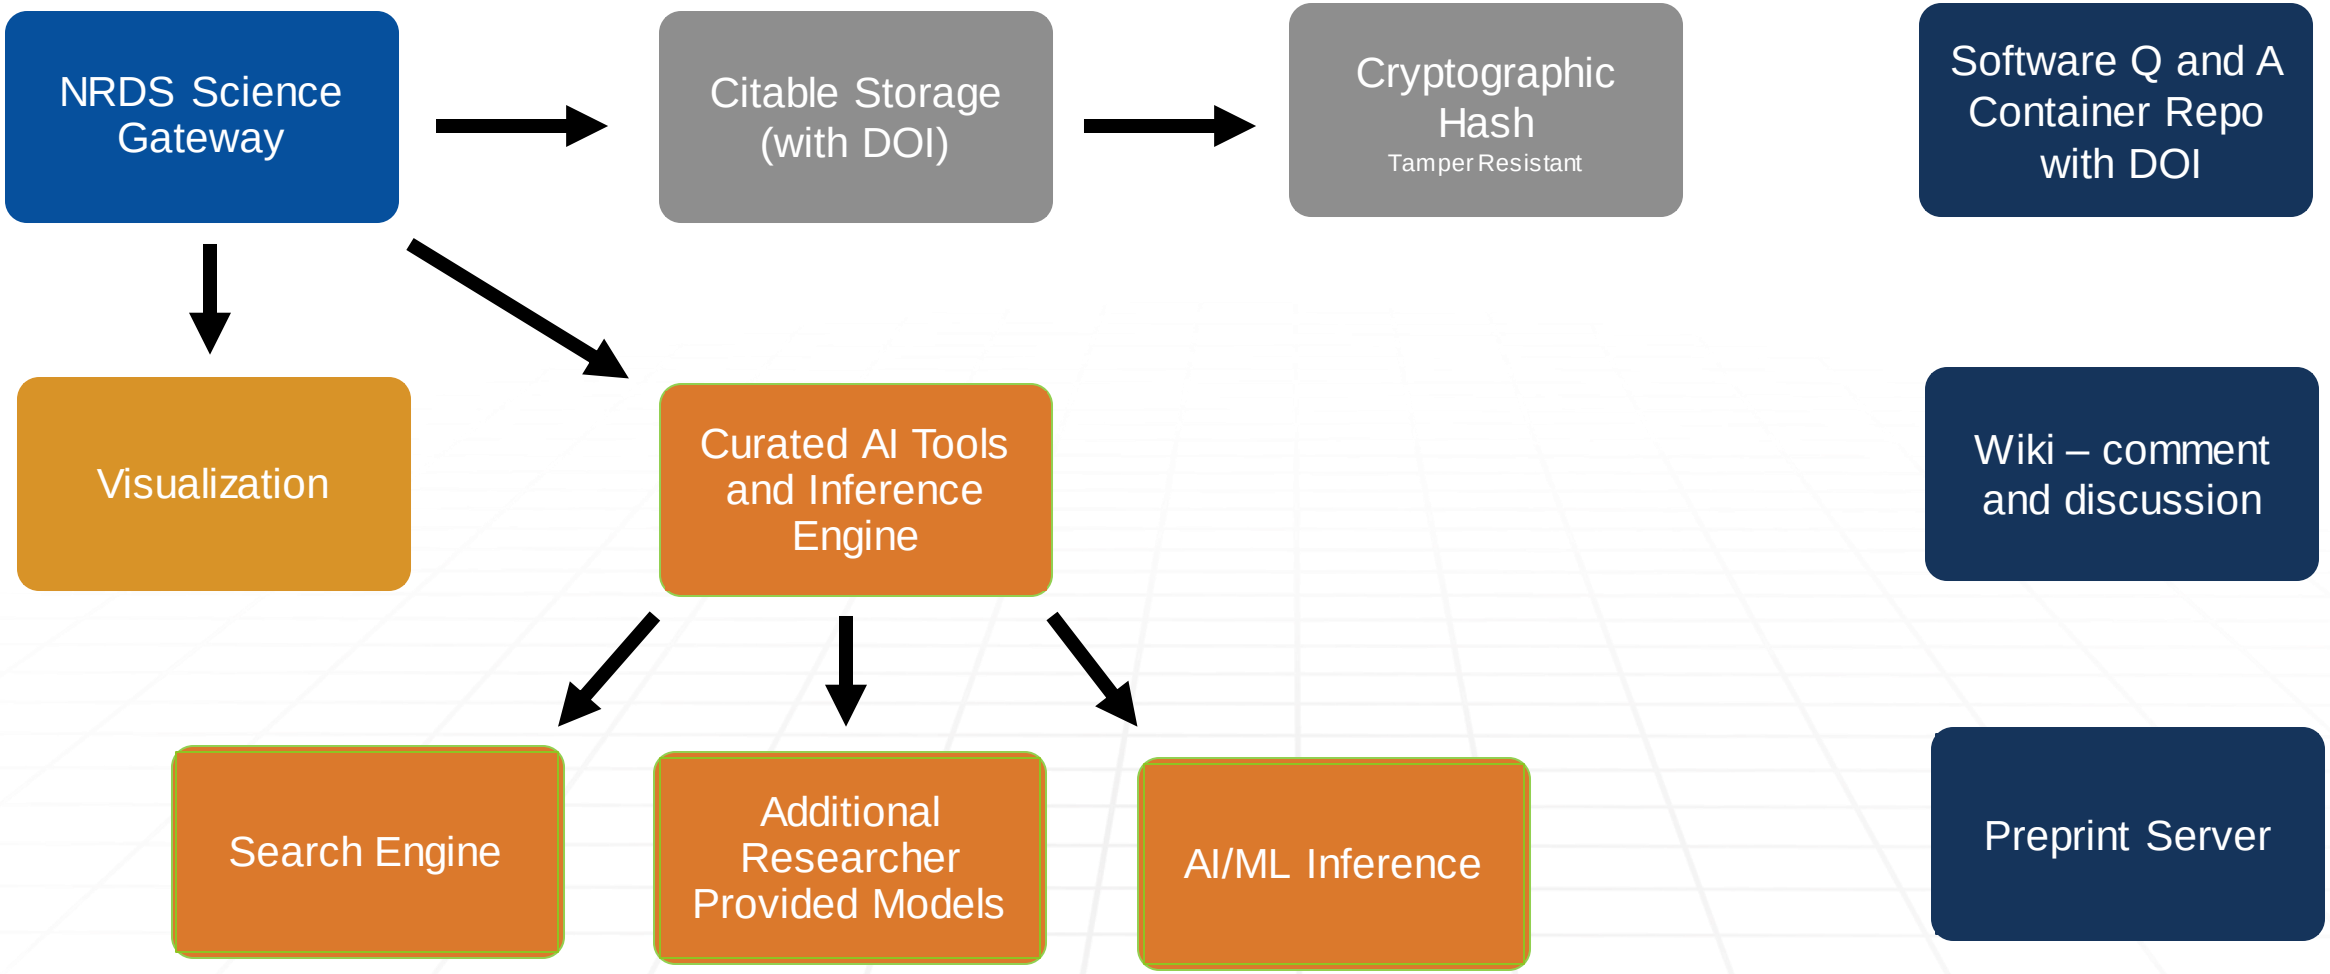
Data Management Strategy – Nuclear Research Data System (NRDS)

The Data Science Gateway

- Centralized location for all nuclear research data
- Data is close to compute
- Allows for downloads and in browser viewing of data
- AI Analysis
- Available for public use in 2024

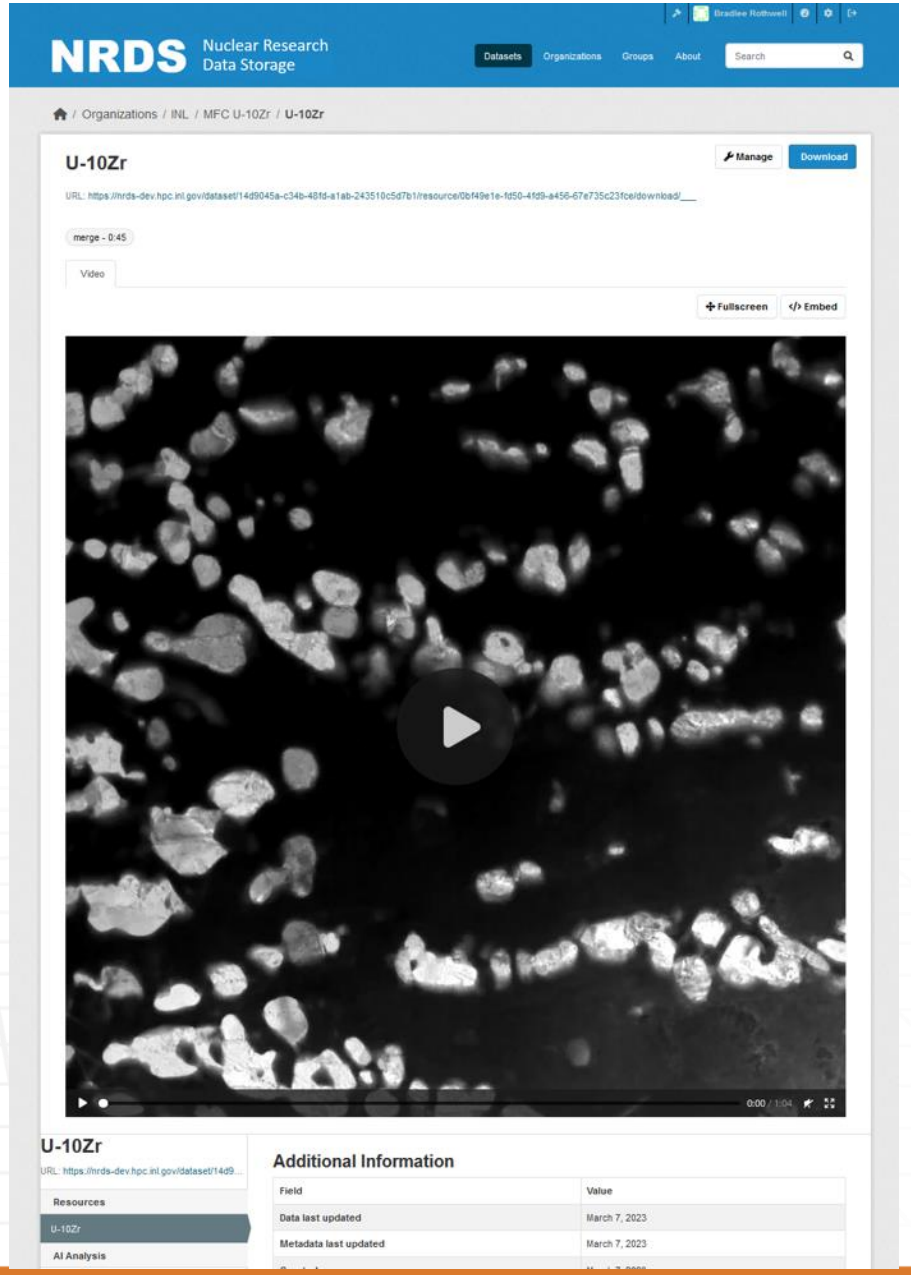


Data System Roadmap

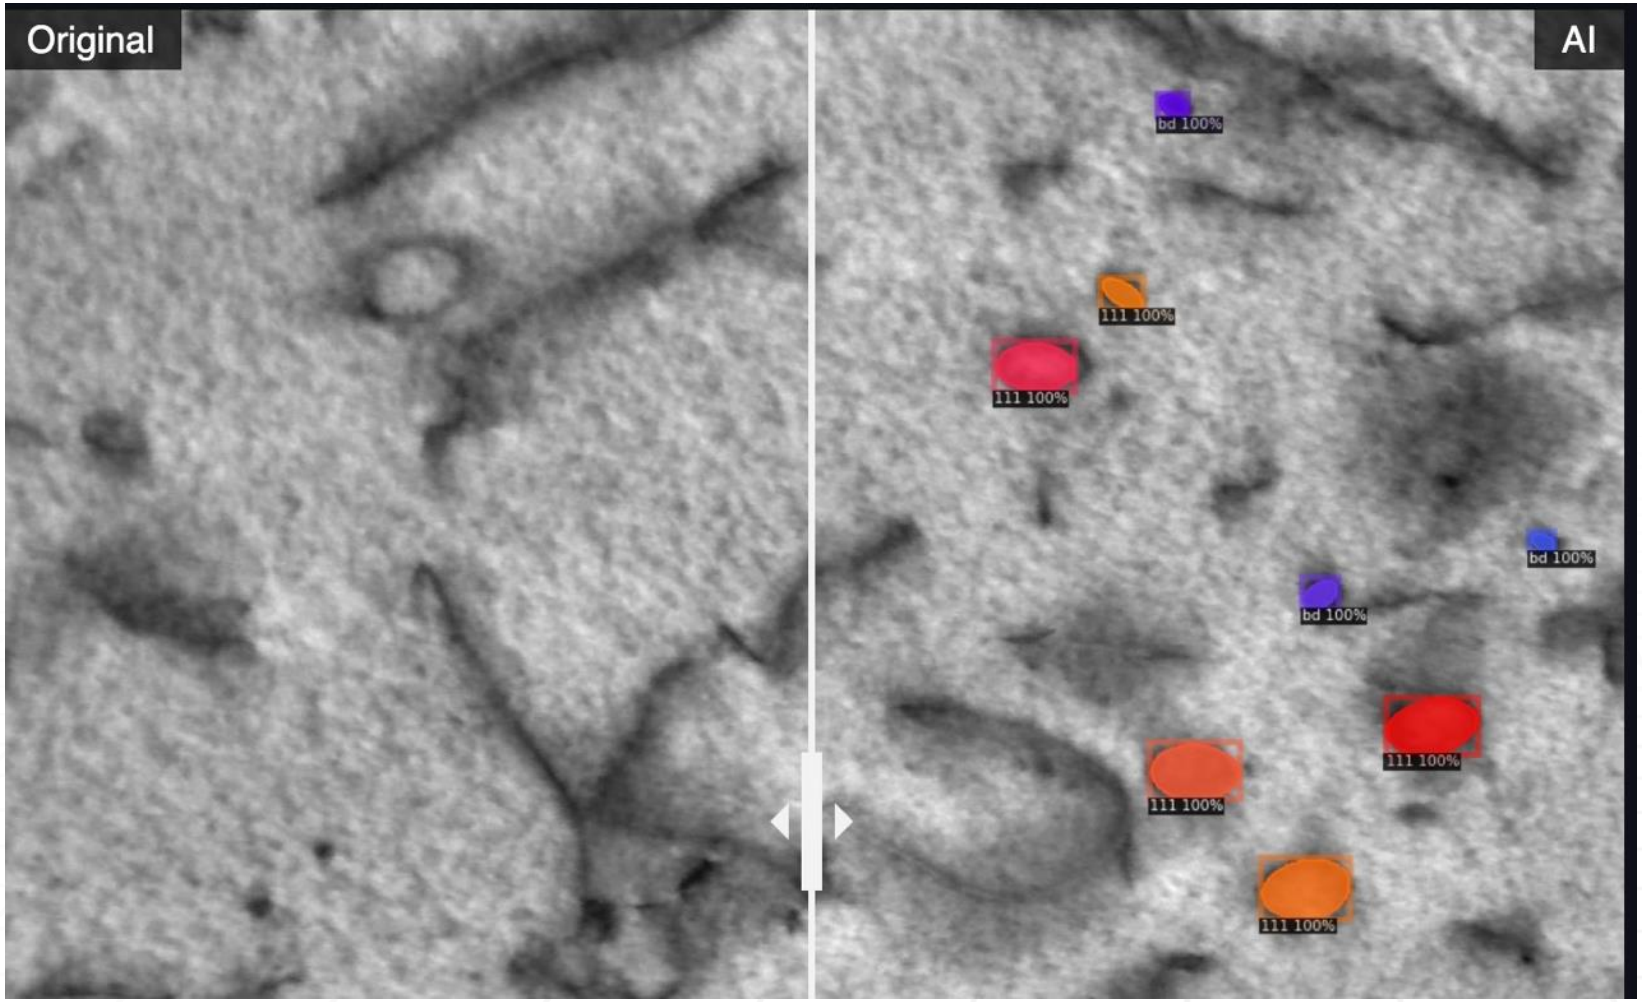


AI Capabilities

- AI Analysis
  - Search Engine
  - Super Resolution
  - Areal Analysis
  - Object Detection
  - Activity Detection
  - Anomaly Detection
    - Regions of Interest
- Visualization
  - 2D Annotation
  - Collaborative 3D Annotation
- Preprint Server
  - Citing Articles
  - NLP Analysis
  - Relevancy Analysis

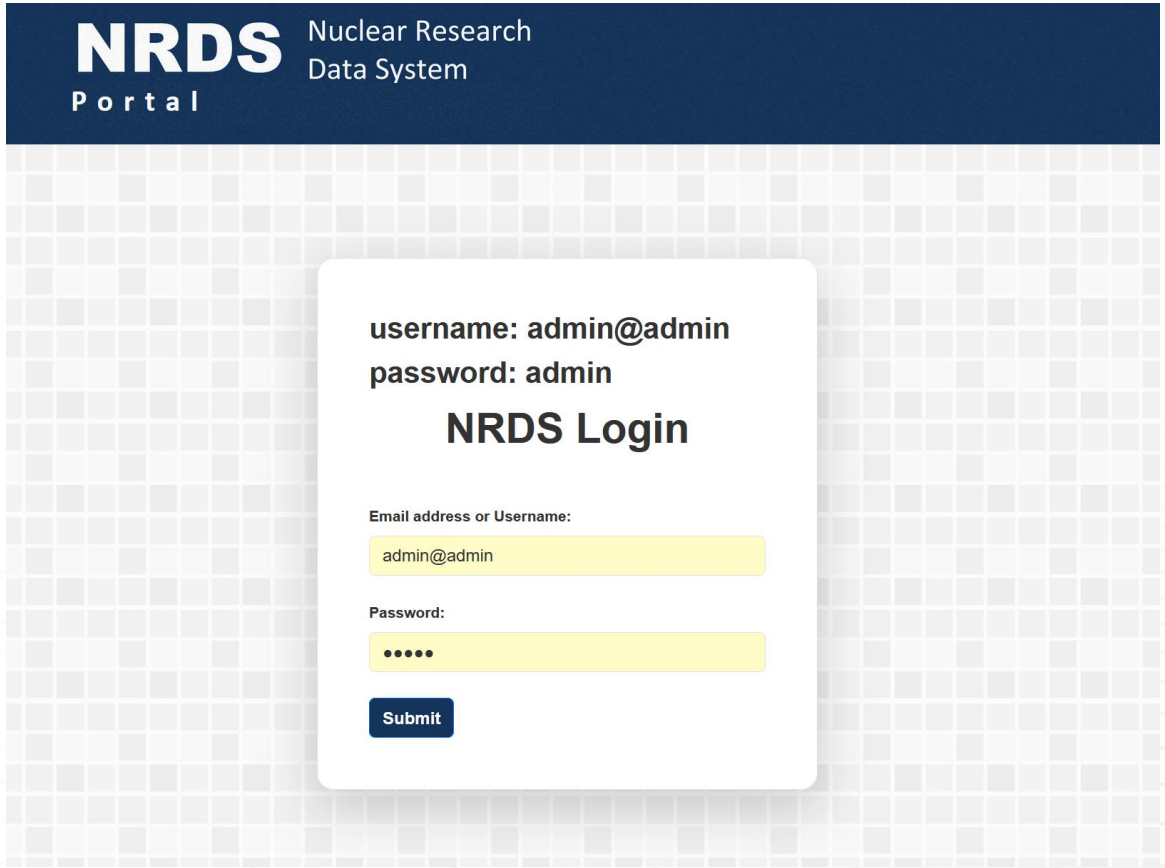


Automatic Image Labeling



Source: Ryan Jacobs, Mingren Shen, Yuhan Liu, Wei Hao, Xiaoshan Li, Ruoyu He, Jacob R.C. Greaves, Donglin Wang, Zeming Xie, Zitong Huang, Chao Wang, Kevin G. Field, Dane Morgan, Performance and limitations of deep learning semantic segmentation of multiple defects in transmission electron micrographs, Cell Reports Physical Science, Volume 3, Issue 5, 2022, 100876, ISSN 2666-3864, https://doi.org/10.1016/j.xcrp.2022.100876.

Data management



Access using NSUF credentials

Data management

Home

Datasets

Create New Dataset

Table Sorting: Title - Dataset Description - Instrument - Tags - License - Organization - Project - Embargo Duration - Owner - Status - Date Created - Date sent Review - Date complete

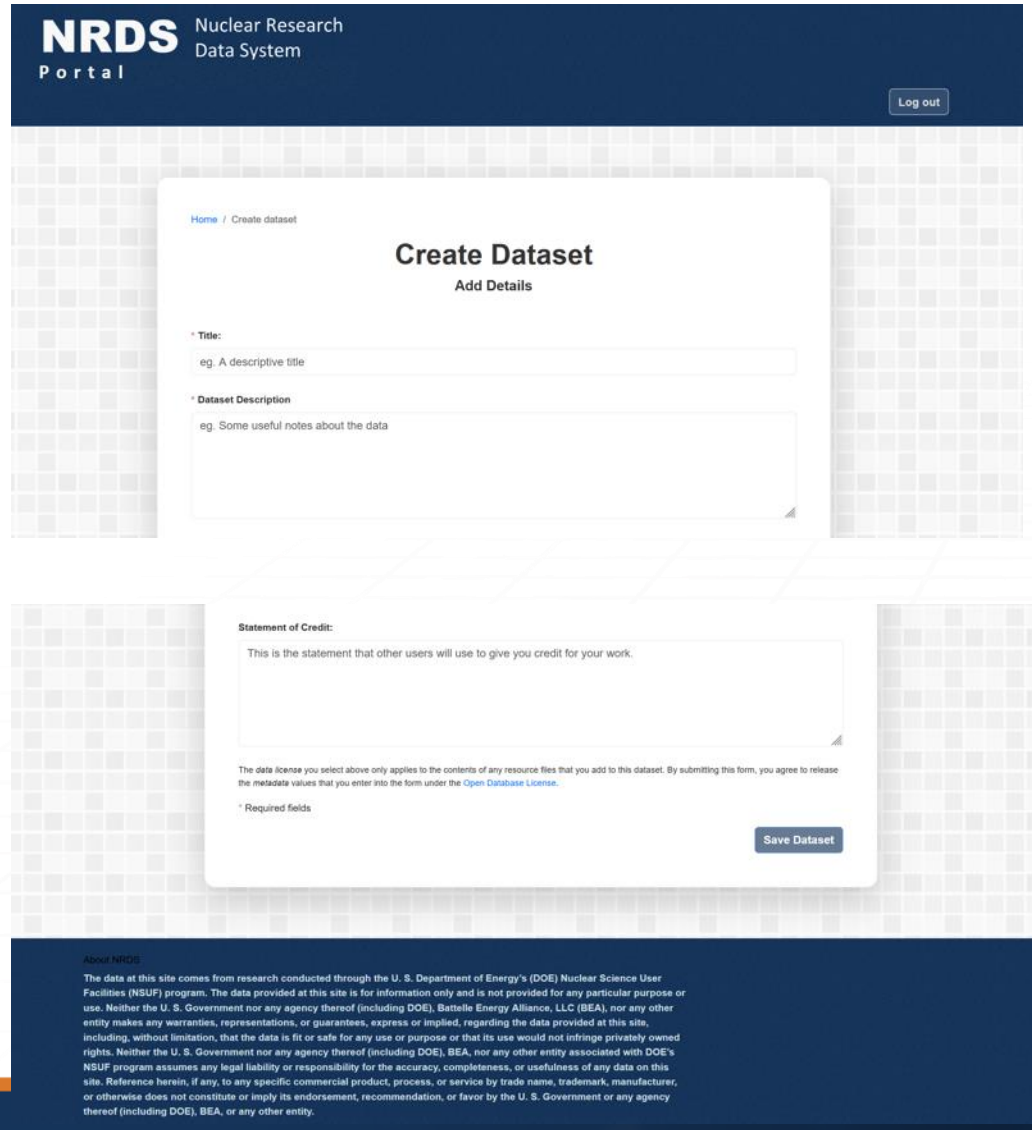
Show 5 **entries** **Search:**

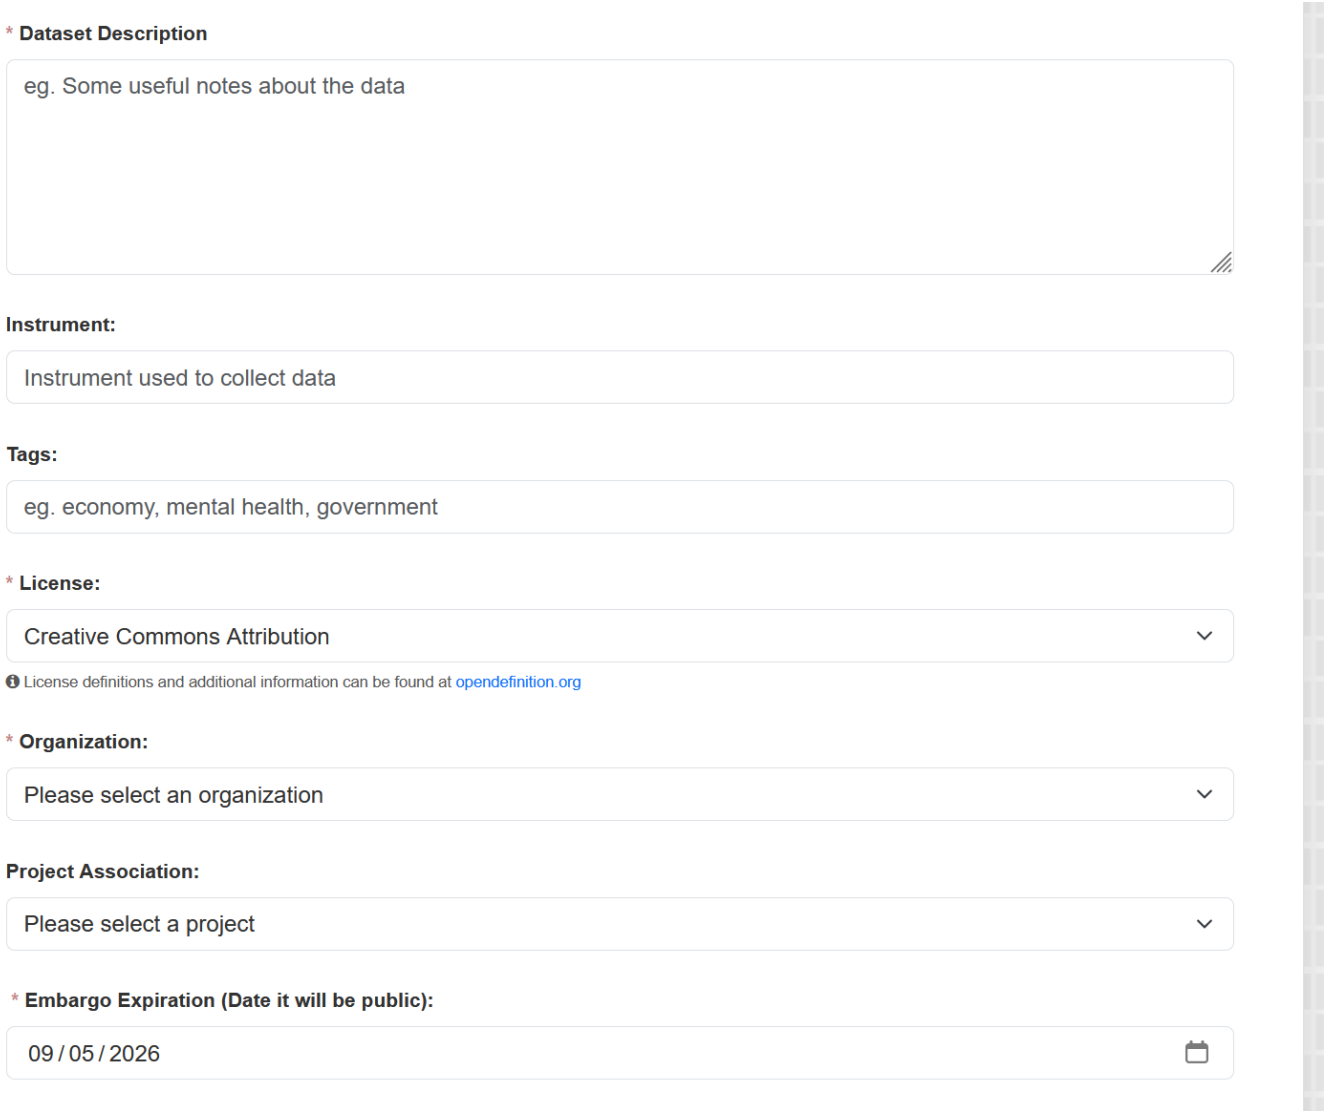
Title	Dataset Description	Instrument	Tags	License	Organization	Project	Embargo Duration	Owner	Status	Date Created	Date sent Review	Date Complete	
Final dataset3	final dataset to upload	---	nuclear	cc-zero	NRDS	NSUF 08-152	2026-06-21	Owen Taylor	pending	2023-08-30	---	---	
Awarded Projects3	List of award projects	---	economy	cc-nc	organization	NSUF 08-154	2025-07-21	NRDS	complete	---	---	---	
Completion Report3	All the completed Reports	---	nuclear	cc-zero	NRDS	NSUF 08-154	2026-07-03	Hoffman	adding	---	---	---	Add data
RTE2	RTE description	---	economy	cc-by-sa	NRDS	NSUF 08-154	2025-05-20	Pittsburgh	pending	---	---	---	
Ohio State2	Datasets from Ohio state	---	economy	cc-nc	NRDS	NSUF 08-153	2025-07-21	Ohio State	adding	---	---	---	Add data
Title	**Dataset Description**	**Instrument**	**Tags**	**License**	**Organization**	**Project**	**Embargo Duration**	**Owner**	**Status**	**Date Created**	**Date sent Review**	**Date Complete**	

Showing 1 to 5 of 23 entries

Previous 1 2 3 4 5 Next

Data management





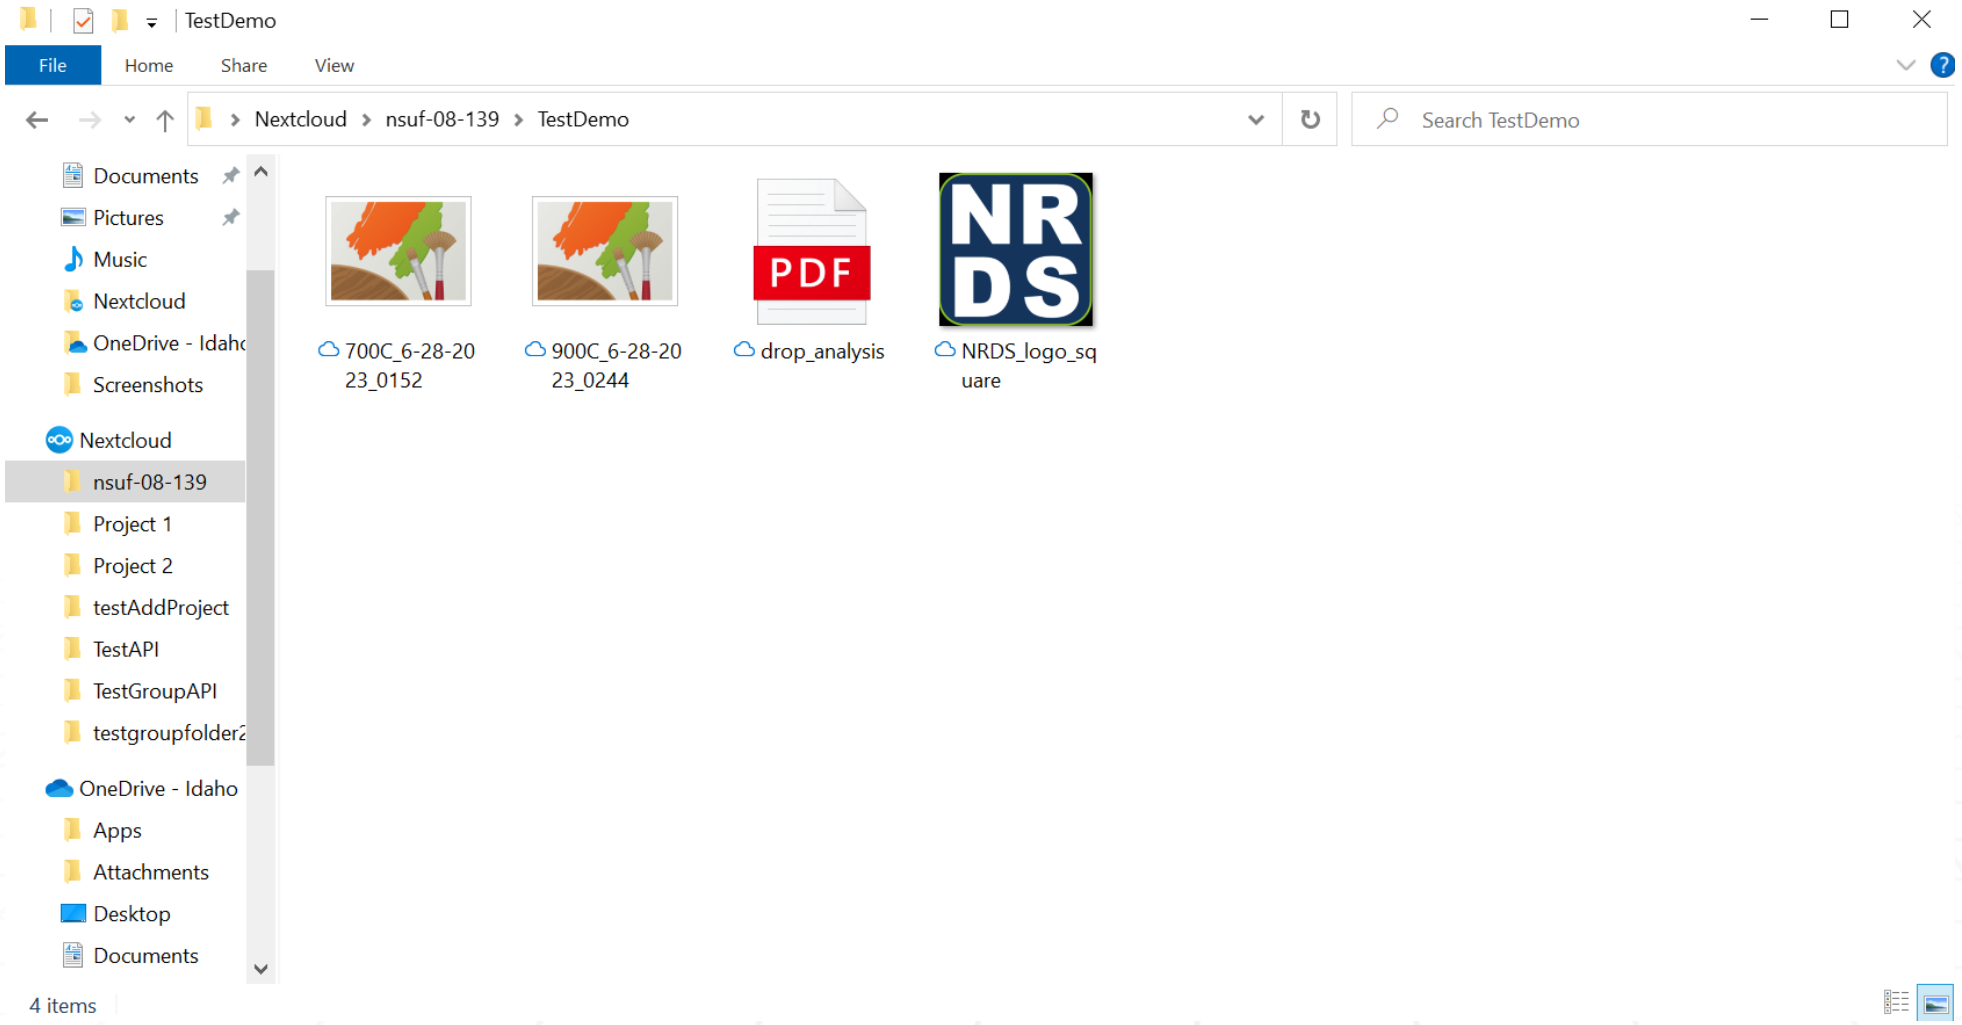
NSUF Nuclear Science User Facilities

Data management

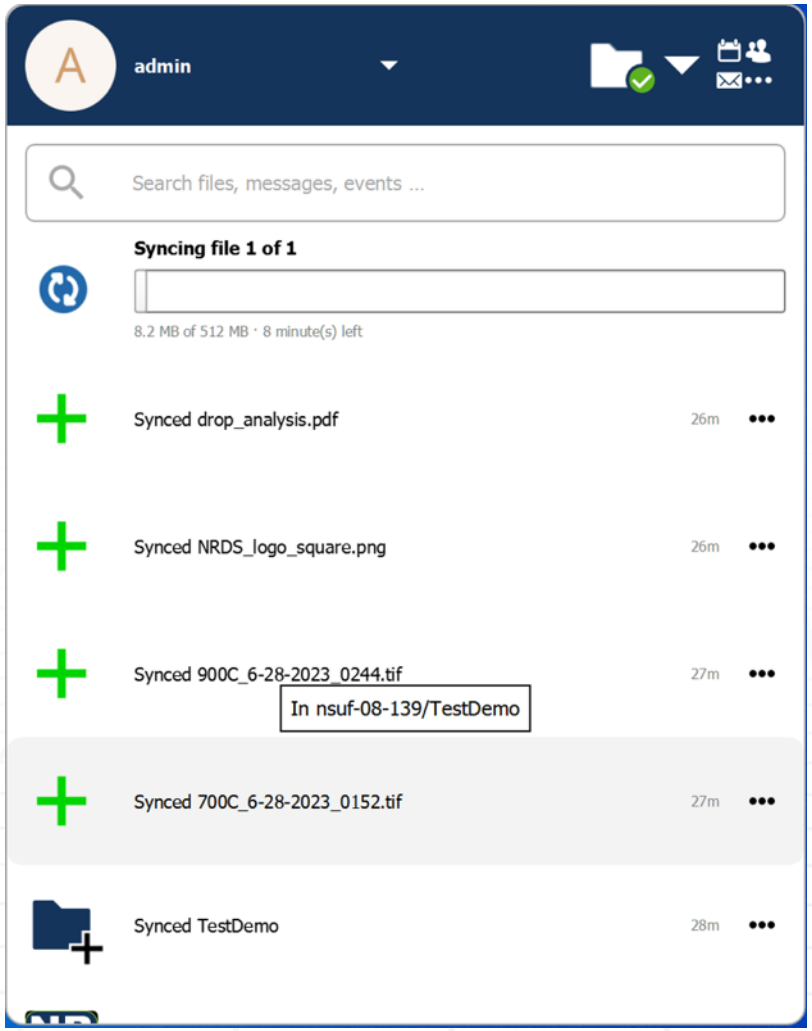


Groups are siloed from different project user members

Uploading/downloading data: desktop/laptop app

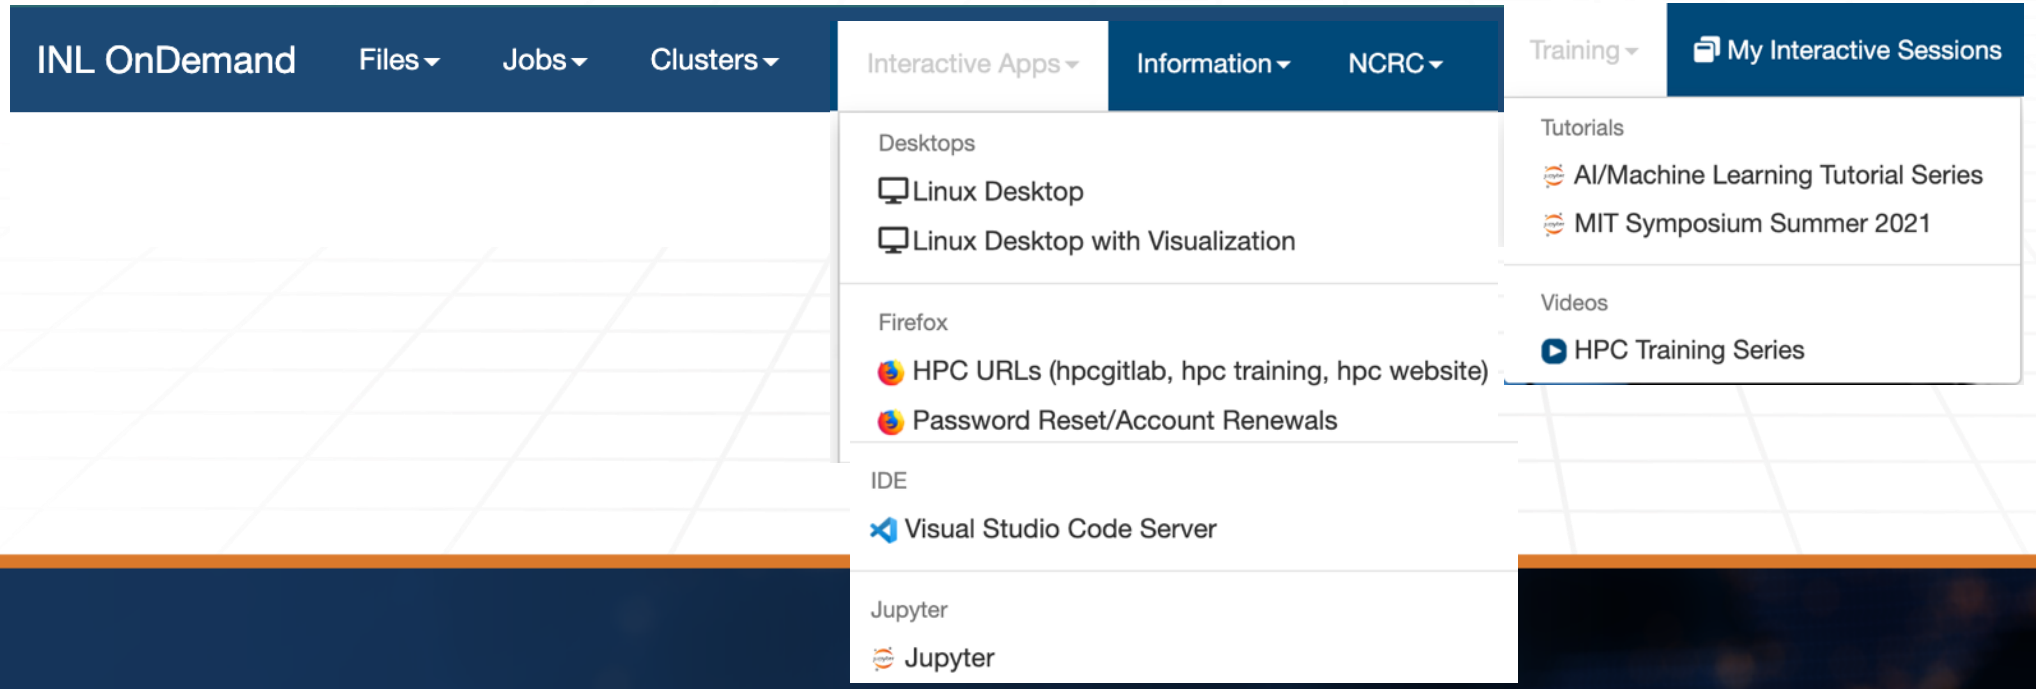


Uploading/downloading data: Desktop/laptop app



HPC Access – Web Portal, Open OnDemand

- OnDemand is a browser based HPC Portal
- Provides HPC filesystem access, SSH access, and login/compute node access
  - Linux Virtual Desktop
  - Jupyter Notebook
  - Visual Studio Code
  - Training and Tutorials



Jupyter Form - CPU

Jupyter version: 780c7c4

This app will launch a Jupyter Lab or Notebook server.

Project

Jupyter Backend

CPU - Sawtooth

Select what type of computational hardware you'd like to have attached to your Jupyter session.

CPUs Requested

1

Cluster	Max Cores
Falcon	36
Lemhi	40
Sawtooth	48

Number of Hours

72

Warning max walltime could be different between systems.
Please see queues secton on the cluster queues for more information.

Backend	Max Hours
CPU - Falcon	336
CPU - Lemhi	72
CPU - Sawtooth	168
GPU - Sawtooth	168
GPU - Hoodoo	168

☐ Use advanced submission settings

Use advanced settings to change your Jupyter server type, number of nodes for your job, or enter project information.

Launch

Jupyter Form - GPU

Jupyter Backend

GPU - Sawtooth

Select what type of computational hardware you'd like to have attached to your Jupyter session.

GPUs Requested

1

Min 1 | Max 4. Requesting GPUs changes the amount of CPUs requested.

Jupyter Form – Advanced Settings

☑ Use advanced submission settings

Use advanced settings to change your Jupyter server type, number of nodes for your job, or enter project information.

Jupyter server type

Lab

Launch the Jupyter session with the Lab or Notebook interface.

Jupyter Working Directory

Enter the directory the Jupyter session will start in

Additional Modulefile Paths

List the directory that contains the modulefiles folder. Multiple directories can be listed (space separated).

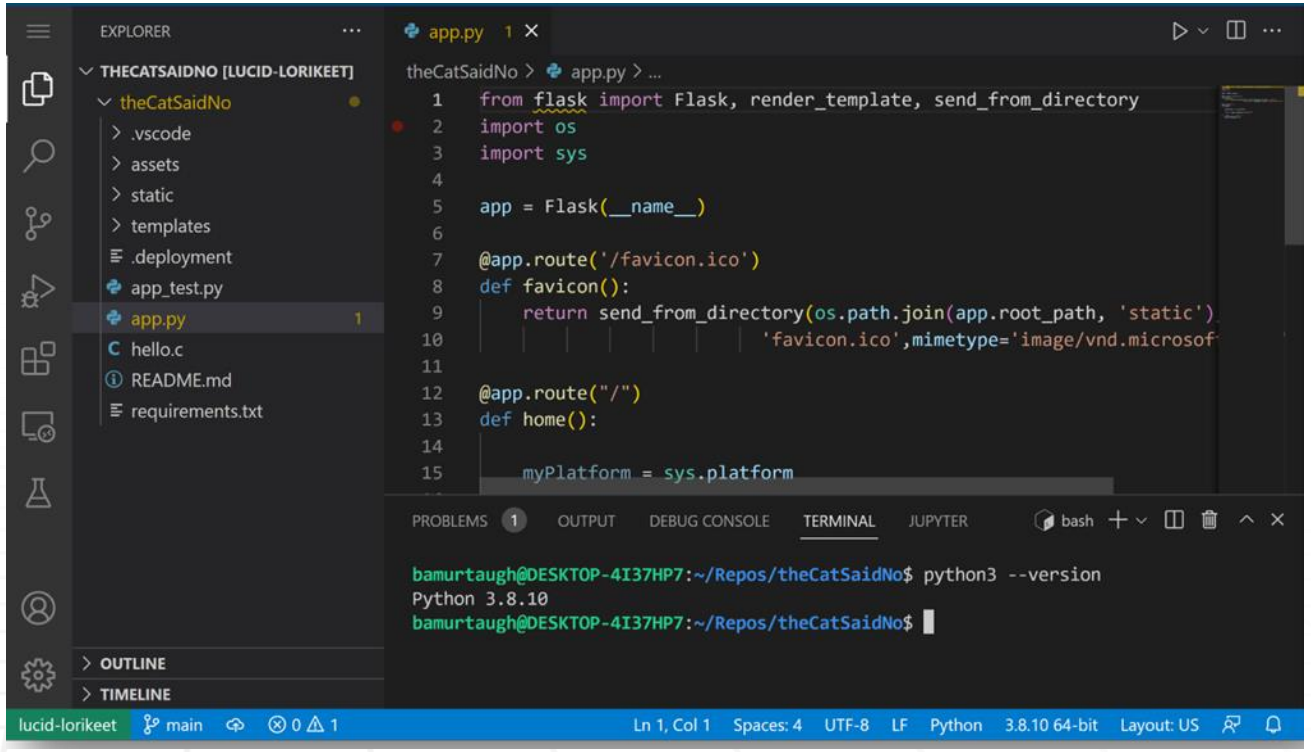
- For example, if the directory was /projects/hpc_test/modulefiles enter **/projects/hpc_test**

Modules

List the module to be loaded. Multiple modules can be listed (space separated). If left blank, the default system python3 will be loaded.

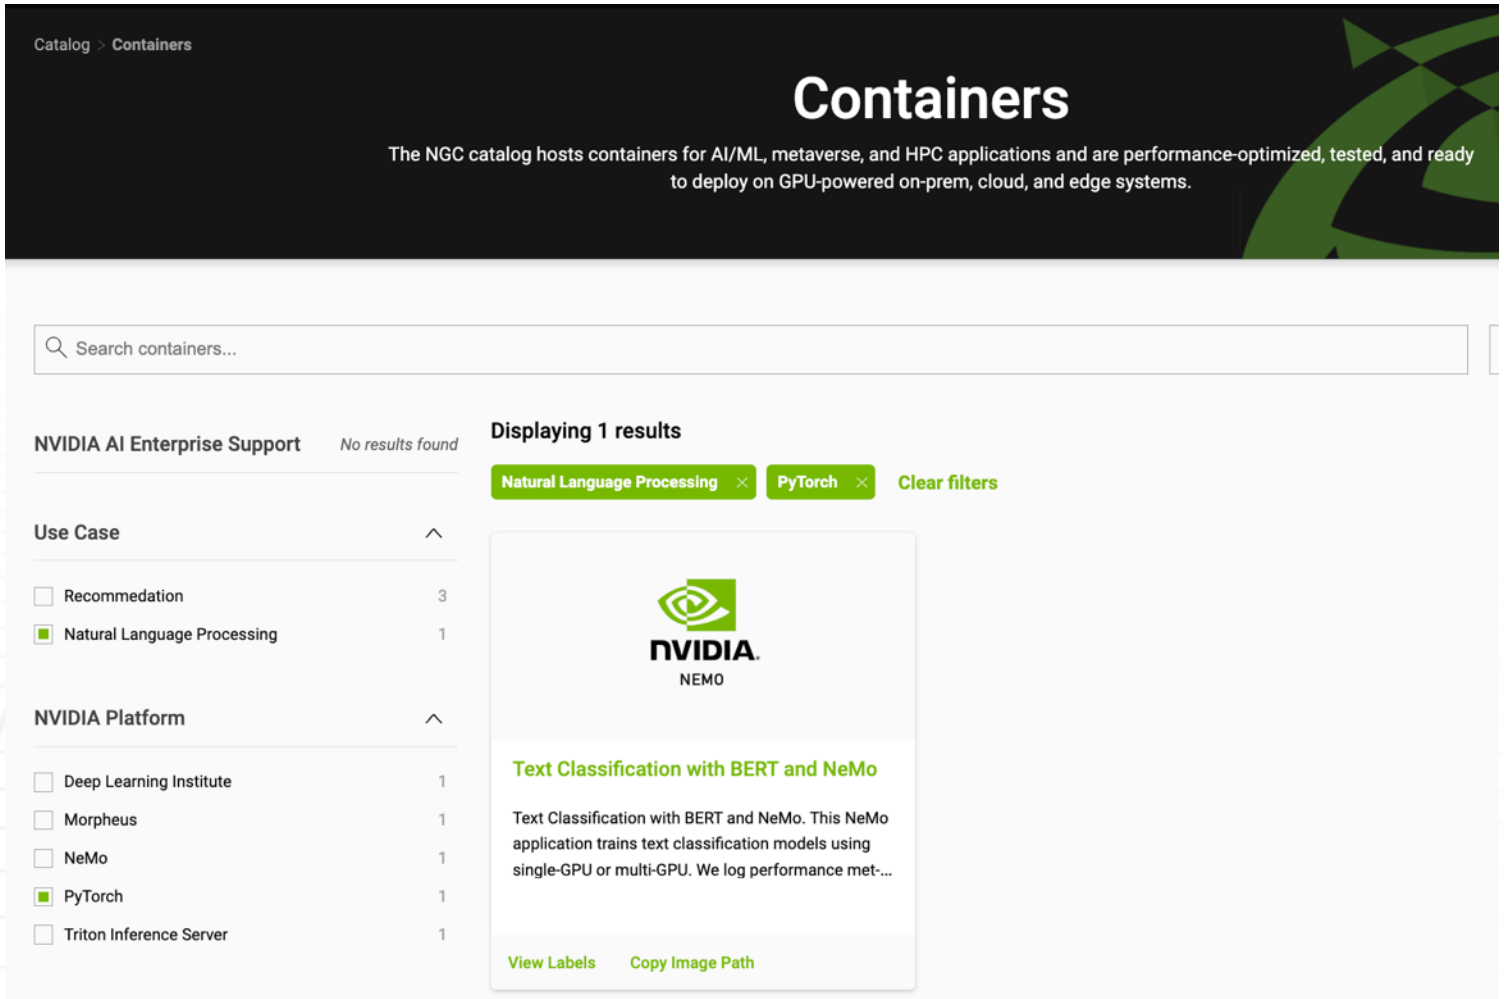
We also have Visual Studio Code Server

- More traditional IDE
- Access your HPC files in an editor from your browser
- Meant for code editing, not for running code. Need to submit additional job for more resources



Bring your own container or let us help

- We convert docker containers to Apptainer so they can run without root
- You can build docker containers locally and we can move them to HPC
- You can also build Apptainer containers on HPC

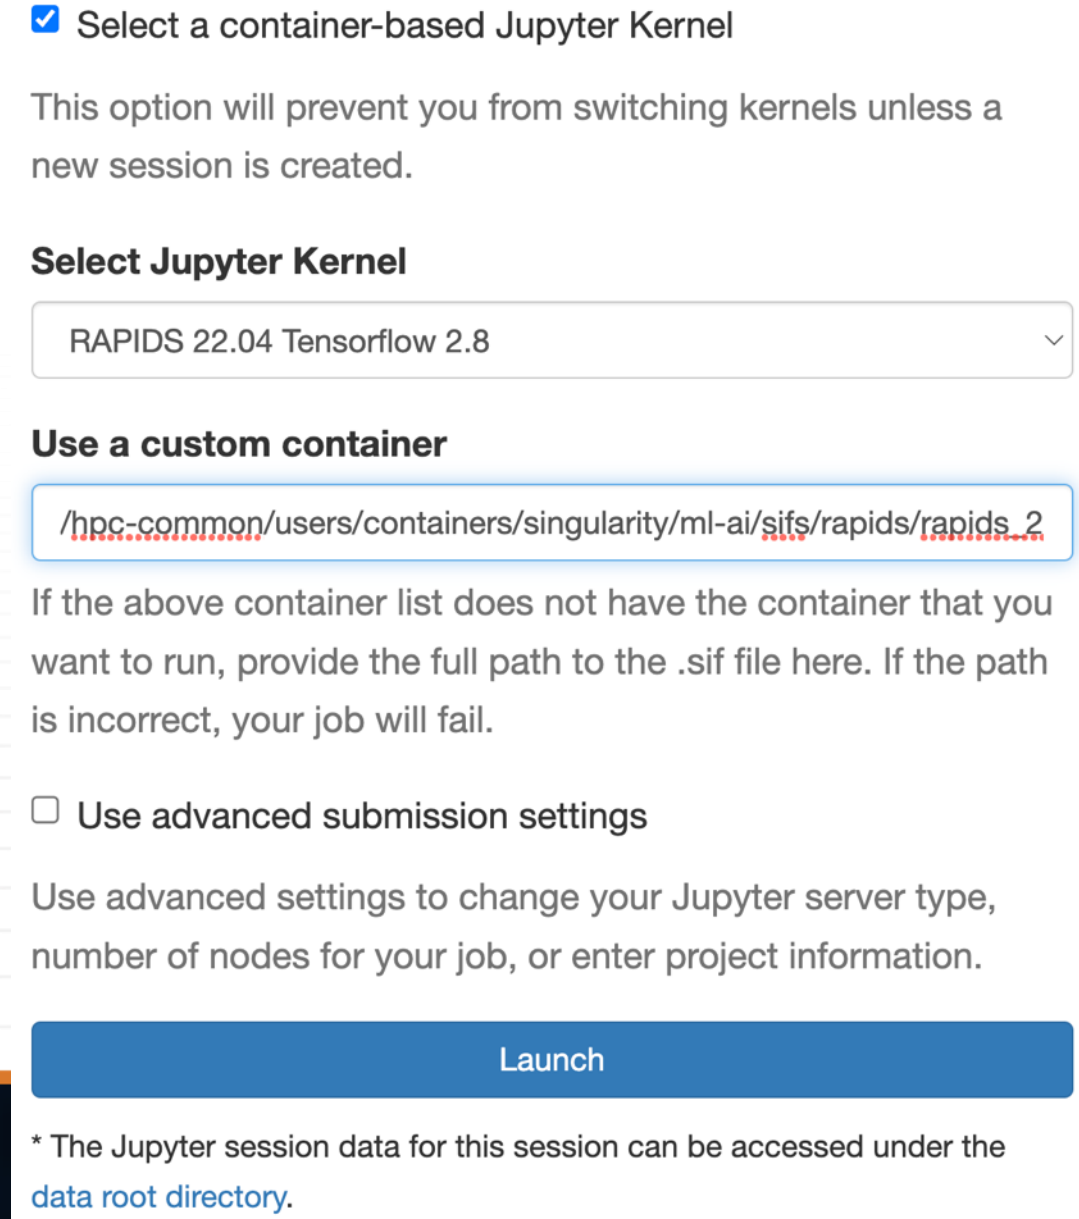


Current Containers

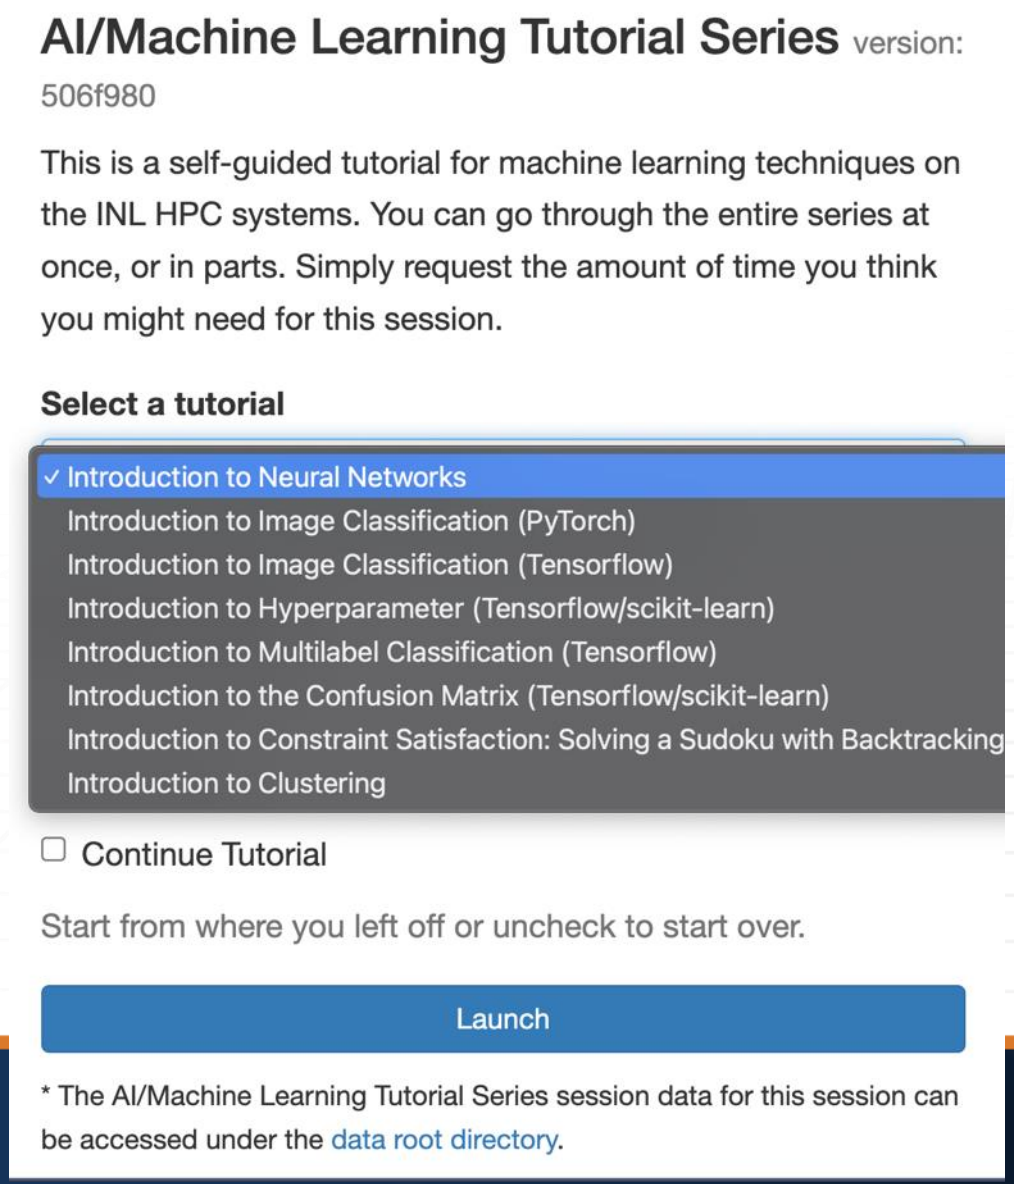
```
/hpc-common/users/containers/singularity/ml-ai/sifs/
├── nvidia
│   ├── nvidia_rapids_22.04_pytorch_1.12.sif
│   ├── nvidia_rapids_22.04_pytorch_1.12_tf_2.8_transformers.sif
│   ├── nvidia_rapids_22.04_pytorch_1.12_transformers_geometric.sif
│   ├── nvidia_rapids_22.04_pytorch_1.12_transformers_geometric-st.sif
│   ├── nvidia_rapids_22.04_pytorch_1.12_transformers.sif
│   ├── nvidia_rapids_22.04_tf_2.8_jax.sif
│   ├── nvidia_rapids_22.04_tf_2.8.sif
│   └── nvidia_tf_2.1.sif
├── rapids
│   ├── rapids_22.04_base.sif
│   └── rapids_22.04_extras.sif
└── tensorflow
    ├── tensorflow_2.8.0_base.sif
    └── tensorflow_2.8.0_extras.sif
```

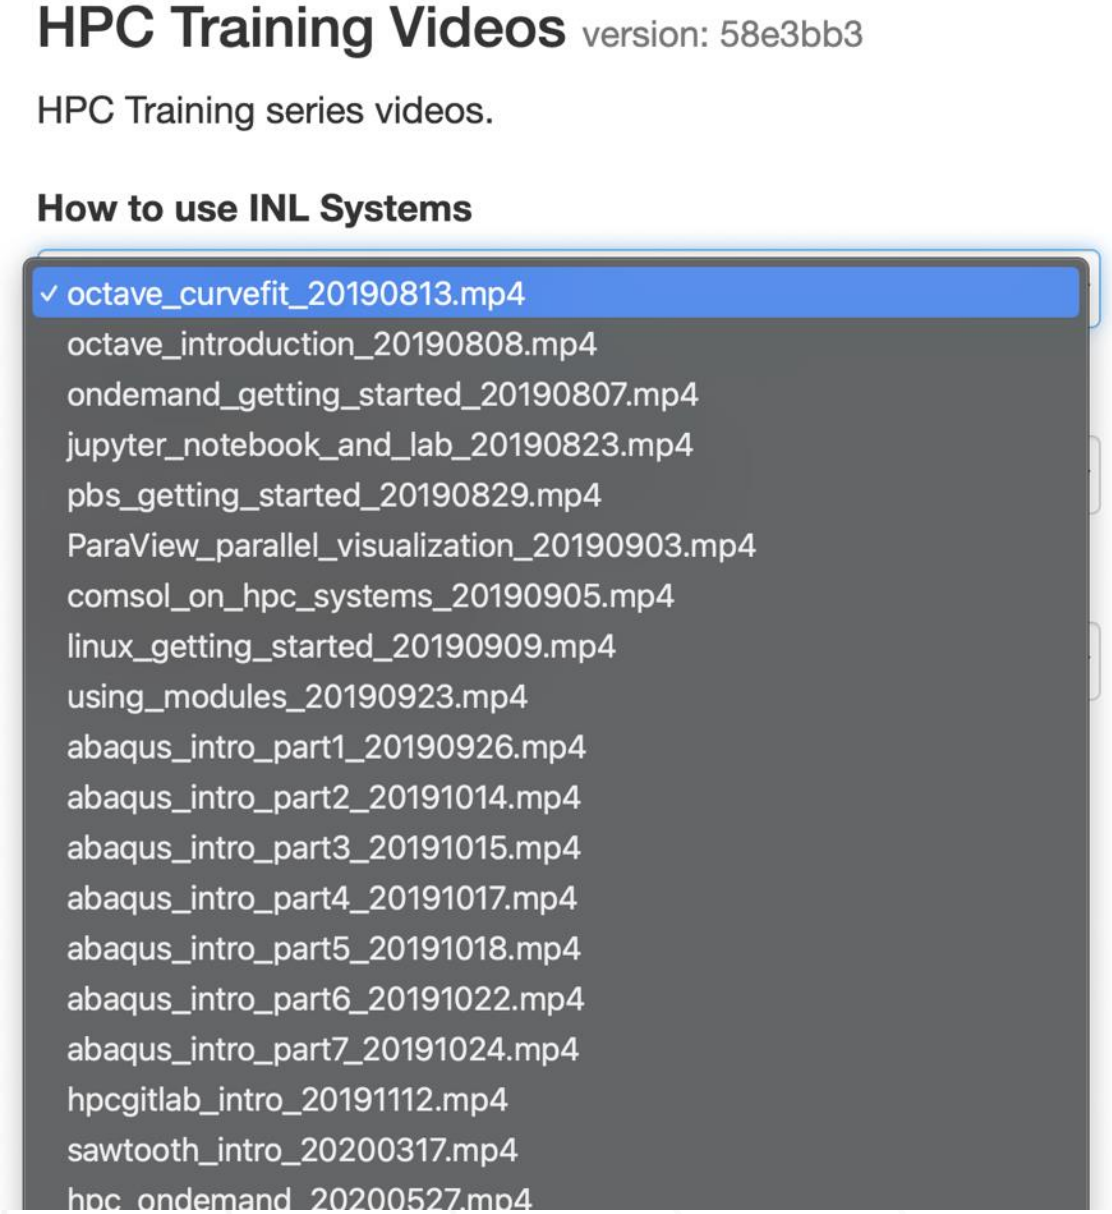
Using Containers in Jupyter Notebooks

- /hpc-common/users/containers/singularity/ml-ai/sifs/rapids/rapids_22.04_base.sif



Training Materials



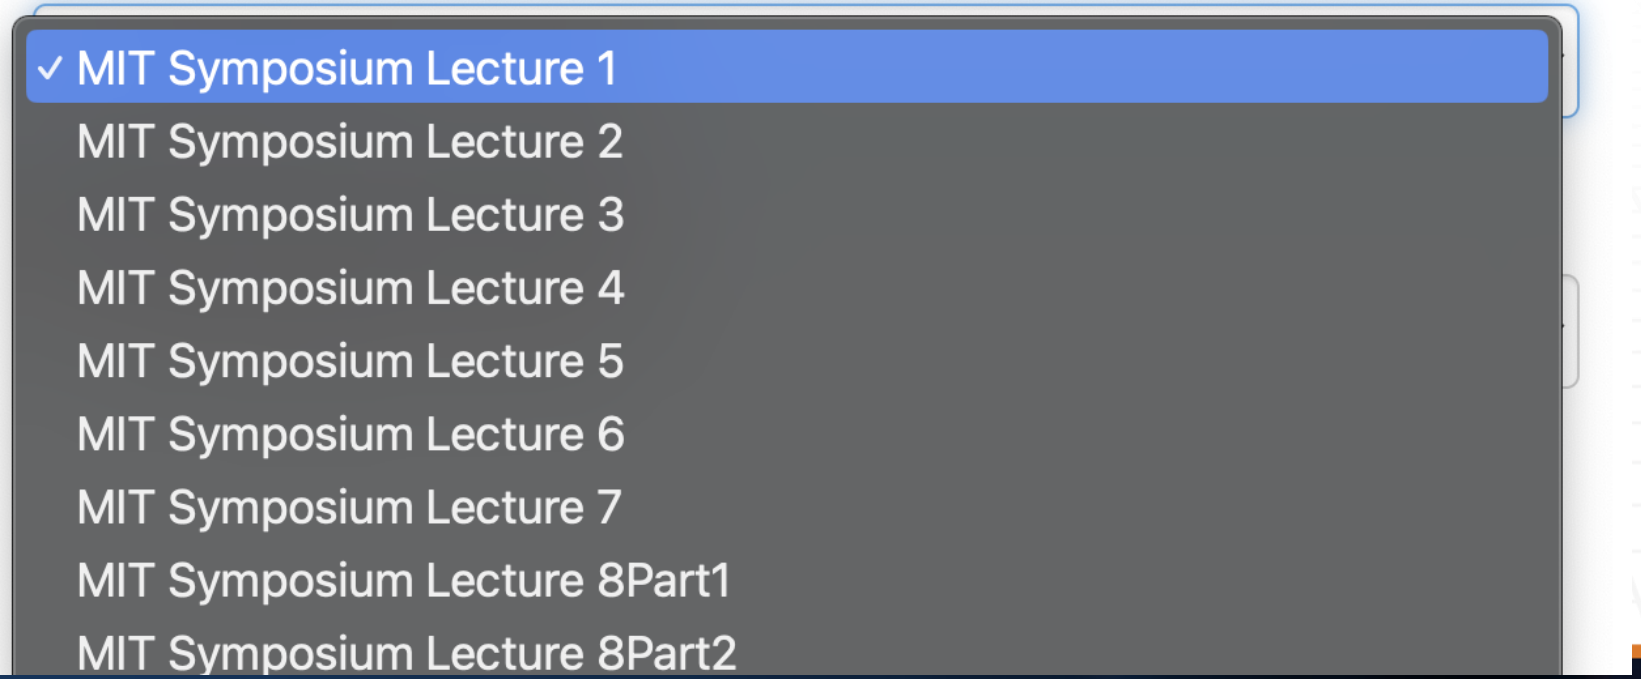


Training Materials

AI/ML Training Videos

AI/ML Training series videos.

AI/ML Symposium 2021



Accessing Systems at INL

- Get an account - https://inl.gov/ncrc
- Command line (SSH) or Web (hpcondemand/ondemand)
- External - hpclogin.inl.gov or hpcondemand.inl.gov (2FA)
- Internal – login node (sawtooth1/2, hoodoo1) or ondemand.hpc.inl.gov
- Submit a job