



Introduction to NCAR HPC Resources

2024 CESM Tutorial

Rory Kelly
HPC Consultant
August 5, 2024

Getting Help

HPC Documentation - <https://ncar-hpc-docs.readthedocs.io>

The screenshot shows the NCAR HPC Documentation website. The header includes the NCAR CISL logo, the text "NCAR HPC Documentation", a search bar, and a GitHub repository link. The main content area features the NCAR Computational & Information Systems Lab logo and the title "User Documentation for NSF NCAR High Performance Computing". A sidebar on the left lists navigation options such as "Home", "Getting Started", "Storage Systems and Data Transfer", "User Environment and HPC Software", "Compute Systems and Services", "Running Jobs (PBS)", "Tutorials", "User Support", "NCAR User Group (NHUG)", and "Contributing to the Documentation". The main text describes the documentation's purpose and provides a list of "Selected Links" including "Getting Started", "Using Derecho", "Using Casper", "Using JupyterHub", and "Getting Help". At the bottom, there is a note about submitting help requests and a contact number.

NCAR CISL NCAR HPC Documentation

GitHub ☆8 👤15

NCAR | COMPUTATIONAL & INFORMATION SYSTEMS LAB

User Documentation for NSF NCAR High Performance Computing

This is the home of the user documentation for the NSF NCAR high-performance computing (HPC) and storage resources managed by CISL. It includes searchable information specific to HPC resources, storage systems, authentication procedures and others, as well as additional how-to articles and troubleshooting articles.

Selected Links

- [Getting Started](#)
- [Using Derecho](#)
- [Using Casper](#)
- [Using JupyterHub](#)
- [Getting Help](#)

Don't find what you need? Log in here to submit a help request: [NCAR Research Computing](#).

You need a CIT password to submit a request. Call **303-497-2400** if you don't have one.



Getting Help

HPC Documentation - <https://ncar-hpc-docs.readthedocs.io>

NCAR CISL | NCAR HPC Documentation

Search

GitHub ☆8 👤15

NCAR HPC Documentation

- Home
- Getting Started >
- Storage Systems and Data Transfer >
- User Environment and HPC Software >
- Compute Systems and Services >
- Running Jobs (PBS) >
- Tutorials >
- User Support
- NCAR User Group (NHUG)
- Contributing to the Documentation

NCAR | COMPUTATIONAL & INFORMATION SYSTEMS LAB

User Documentation for NSF NCAR High Performance Computing

This is the home of the user documentation for the NSF NCAR high-performance computing (HPC) and storage resources managed by CISL. It includes searchable information specific to HPC resources, storage systems, authentication procedures and others, as well as additional how-to articles and troubleshooting articles.

Selected Links

- [Getting Started](#)
- [Using Derecho](#)
- [Using Casper](#)
- [Using JupyterHub](#)
- [Getting Help](#)

Don't find what you need? Log in here to submit a help request: [NCAR Research Computing](#).

You need a CIT password to submit a request. Call **303-497-2400** if you don't have one.

Page Outline
Selected Links

Consultant on Duty



Getting Help

HPC Documentation - <https://ncar-hpc-docs.readthedocs.io>

The screenshot shows the NCAR HPC Documentation website. The header includes the NCAR CISL logo, the text "NCAR HPC Documentation", a search bar, and a GitHub link. The left sidebar lists navigation categories such as "Getting Started", "Storage Systems and Data Transfer", "User Environment and HPC Software", "Compute Systems and Services", "Running Jobs (PBS)", "Tutorials", "User Support", and "NCAR User Group (NHUG)". The main content area features the NCAR Computational & Information Systems Lab logo and the title "User Documentation for NSF NCAR High Performance Computing". Below the title is a paragraph describing the documentation's scope. A "Selected Links" section lists: "Getting Started", "Using Derecho", "Using Casper", "Using JupyterHub", and "Getting Help". A red arrow points from the text "Consultant on Duty" to the "Getting Help" link. Another red arrow points from the text "Submit a ticket" to the link "NCAR Research Computing" in the text below. At the bottom, there is a note about submitting requests and a phone number: "You need a CIT password to submit a request. Call 303-497-2400 if you don't have one."



Getting Help

HPC Documentation - <https://ncar-hpc-docs.readthedocs.io>

The screenshot shows the NCAR HPC Documentation website. The header includes the NCAR CISL logo, the title 'NCAR HPC Documentation', a search bar, and a GitHub link. The main content area features the NCAR Computational & Information Systems Lab logo and the title 'User Documentation for NSF NCAR High Performance Computing'. A paragraph describes the site's purpose. A 'Selected Links' section lists several articles, with red arrows pointing to 'Getting Help', 'Submit a ticket', and 'Password Help'. A footer section provides information on how to submit a help request, including a phone number.

NCAR CISL NCAR HPC Documentation

NCAR | COMPUTATIONAL & INFORMATION SYSTEMS LAB

User Documentation for NSF NCAR High Performance Computing

This is the home of the user documentation for the NSF NCAR high-performance computing (HPC) and storage resources managed by CISL. It includes searchable information specific to HPC resources, storage systems, authentication procedures and others, as well as additional how-to articles and troubleshooting articles.

Selected Links

- [Getting Started](#)
- [Using Derecho](#)
- [Using Casper](#)
- [Using JupyterHub](#)
- [Getting Help](#)

Don't find what you need? Log in here to submit a help request: [NCAR Research Computing](#).

You need a CIT password to submit a request. Call **303-497-2400** if you don't have one.

Consultant on Duty

Submit a ticket

Password Help



Topics to Cover

- Available systems and their uses
- Signing in and managing data
- Accessing software
- Managing jobs using Batch schedulers
- Customizing your user environment

This is only an introduction; for full documentation, see:

<https://ncar-hpc-docs.readthedocs.io/en/latest/>

Derecho - Primary HPC Resource

- HPC Cray EX, 19.87 petaflops
- 2488 CPU Nodes
 - Dual socket, 64-core AMD 7763 “Milan” CPUs
 - 256 GB DDR4 Memory
 - 1 Cassini Slingshot-11 NIC
- 82 GPU Nodes
 - Single socket, 64-core AMD 7763 “Milan” CPUs
 - 512 GB DDR4 Memory
 - 4 NVIDIA A100 GPUs
 - NVLink GPU interconnect
 - 4 Cassini Slingshot-11 NICs



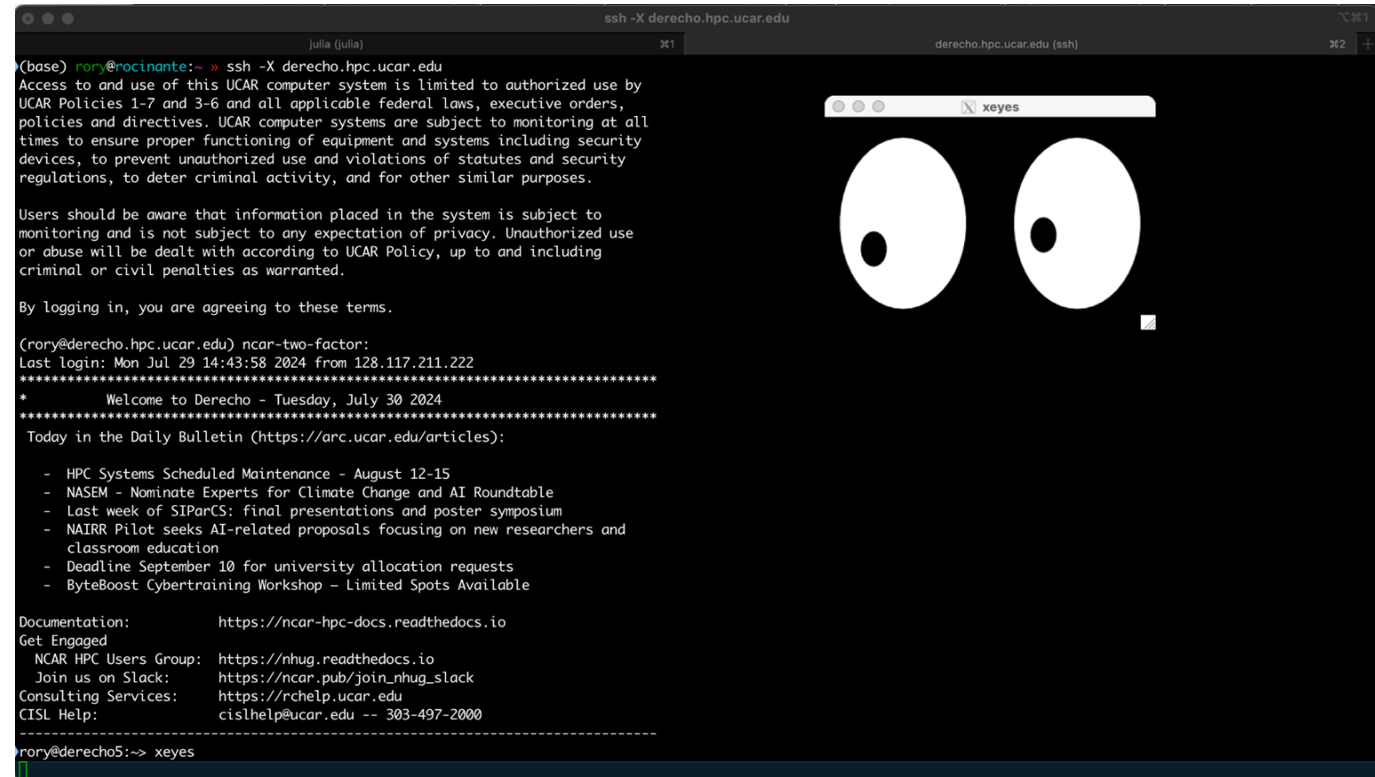
Casper - Data Analysis, HTC, Visualization

- 22 nodes for data analysis and visualization jobs. Each node has 36 cores and up to 384 GB memory.
 - 9 of these nodes also feature an NVIDIA Quadro GP 100 GPU.
- 18 nodes feature large-memory, dense GPU configurations to support explorations in machine learning (ML) and GPU computing
 - 4 of these nodes feature 4 NVIDIA Tesla V100 GPUs
 - 6 of these nodes feature 8 NVIDIA Tesla V100 GPUs
 - 8 of these nodes feature 4 NVIDIA Ampere A100 GPUs
- 64 high-throughput computing (HTC) nodes for small computing tasks using 1 or 2 CPUs.
 - 62 HTC nodes have 384 GB of available memory
 - 2 HTC nodes have 1.5 TB of available memory
- 4 nodes are reserved for Research Data Archive workflows



System Access - Logging in from a Terminal

- Use ssh along with your username to log in
`ssh -XY -l username derecho.hpc.ucar.edu`
`ssh -XY -l username casper.hpc.ucar.edu`
- Use Duo for authentication
- Derecho - 8 login nodes
- Casper - 2 login nodes



```
(base) rory@rocinante:~$ ssh -X derecho.hpc.ucar.edu
Access to and use of this UCAR computer system is limited to authorized use by
UCAR Policies 1-7 and 3-6 and all applicable federal laws, executive orders,
policies and directives. UCAR computer systems are subject to monitoring at all
times to ensure proper functioning of equipment and systems including security
devices, to prevent unauthorized use and violations of statutes and security
regulations, to deter criminal activity, and for other similar purposes.

Users should be aware that information placed in the system is subject to
monitoring and is not subject to any expectation of privacy. Unauthorized use
or abuse will be dealt with according to UCAR Policy, up to and including
criminal or civil penalties as warranted.

By logging in, you are agreeing to these terms.

(rory@derecho.hpc.ucar.edu) ncar-two-factor:
Last login: Mon Jul 29 14:43:58 2024 from 128.117.211.222
*****
*           Welcome to Derecho - Tuesday, July 30 2024           *
*****
Today in the Daily Bulletin (https://arc.ucar.edu/articles):

- HPC Systems Scheduled Maintenance - August 12-15
- NASEM - Nominate Experts for Climate Change and AI Roundtable
- Last week of SIPaCS: final presentations and poster symposium
- NAIRR Pilot seeks AI-related proposals focusing on new researchers and
  classroom education
- Deadline September 10 for university allocation requests
- ByteBoost Cybertraining Workshop - Limited Spots Available

Documentation:           https://ncar-hpc-docs.readthedocs.io
Get Engaged
  NCAR HPC Users Group: https://nhug.readthedocs.io
  Join us on Slack:      https://ncar.pub/join\_nhug\_slack
Consulting Services:    https://rhelp.ucar.edu
CISL Help:              cislhelp@ucar.edu -- 303-497-2000
-----
rory@derecho5:~$ xeyes
```

Run GUI Programs with VNC

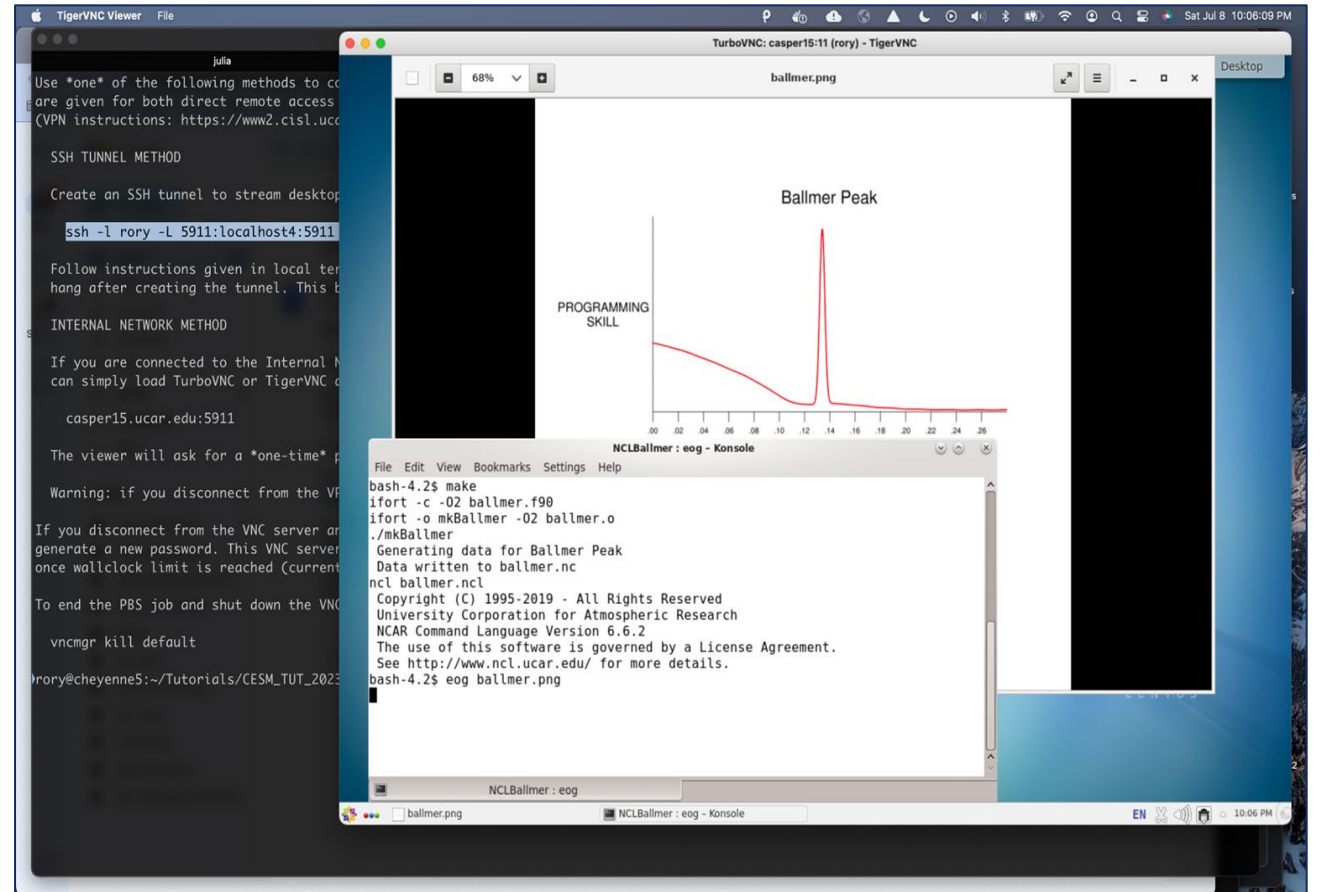
VNC can be used to run a remote GNOME/KDE desktop

Need to install a VNC client first - We recommend TigerVNC, but other VNC clients such as TurboVNC will also work

Usage:

```
vncmgr create -A <project code>
```

```
vncmgr (interactive)
```



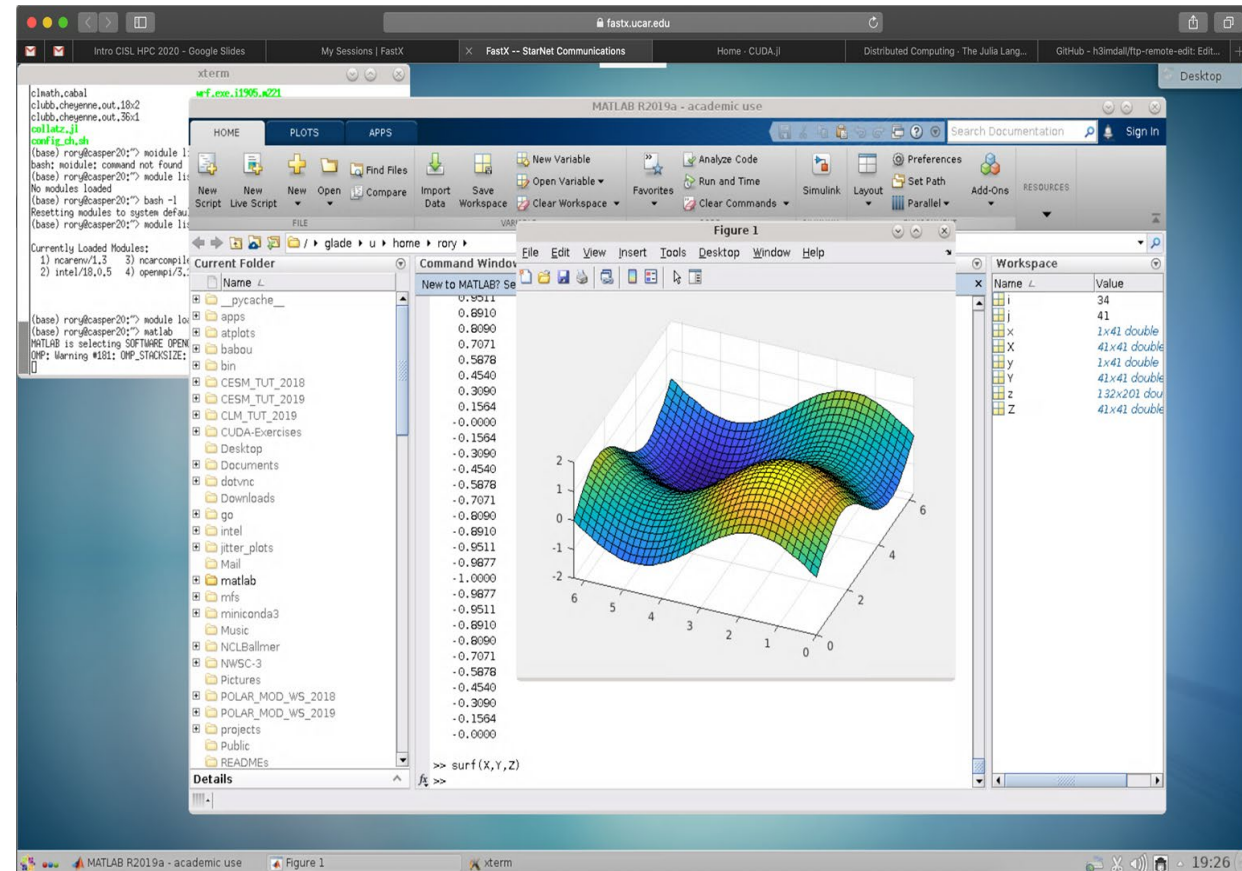
Run GUI Programs with FastX

FastX can be used to run a remote desktop or terminal session in a browser without a client.

Connect to the NCAR VPN, then go to <https://fastx.ucar.edu:3300>

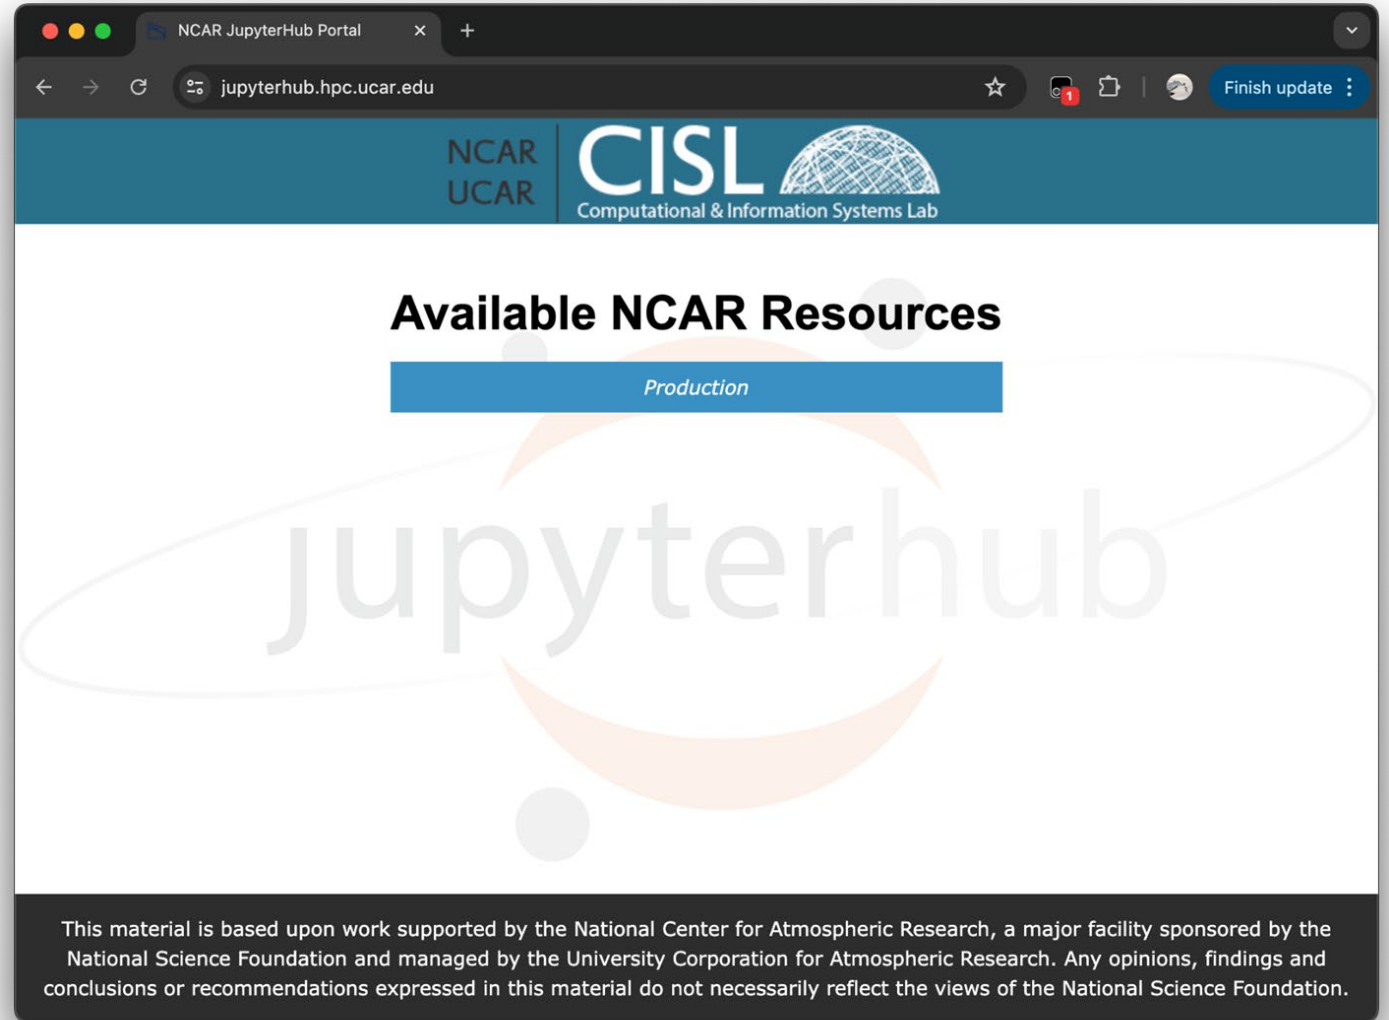
Can also be setup using an SSH tunnel

- `ssh -L 3300:fastx.ucar.edu:3300 username@fastx.ucar.edu` (duo auth)
- Go to <https://localhost:3300> in a browser
 - ignore unsafe warnings
 - duo auth again
- Open a KDE Desktop



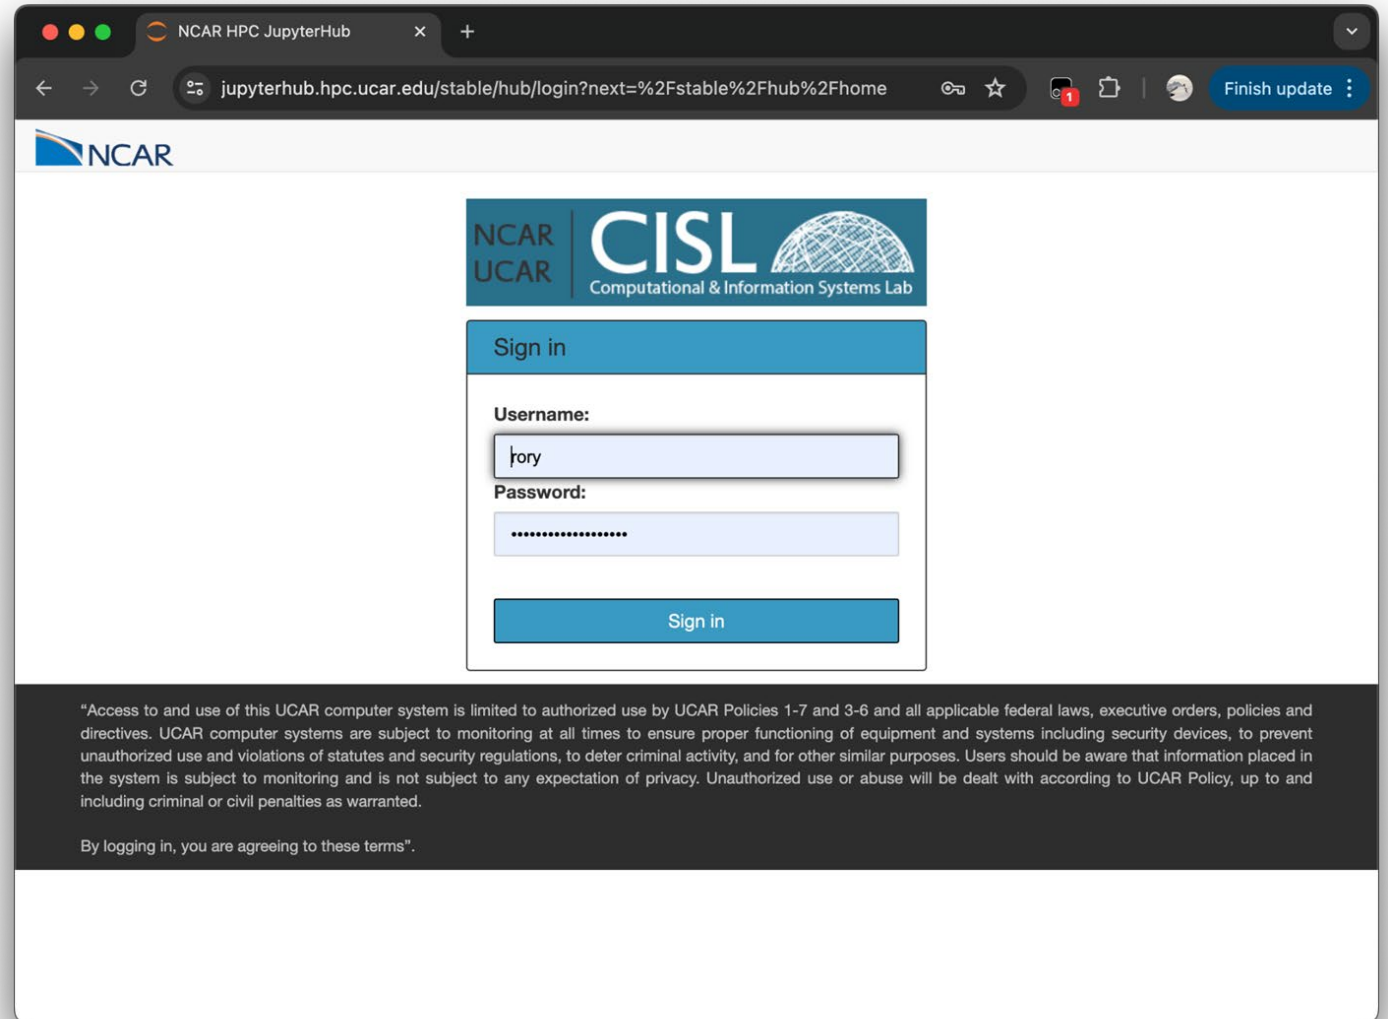
Jupyterhub - Logging In

Derecho and Casper are also accessible via the JupyterHub service at jupyterhub.hpc.ucar.edu



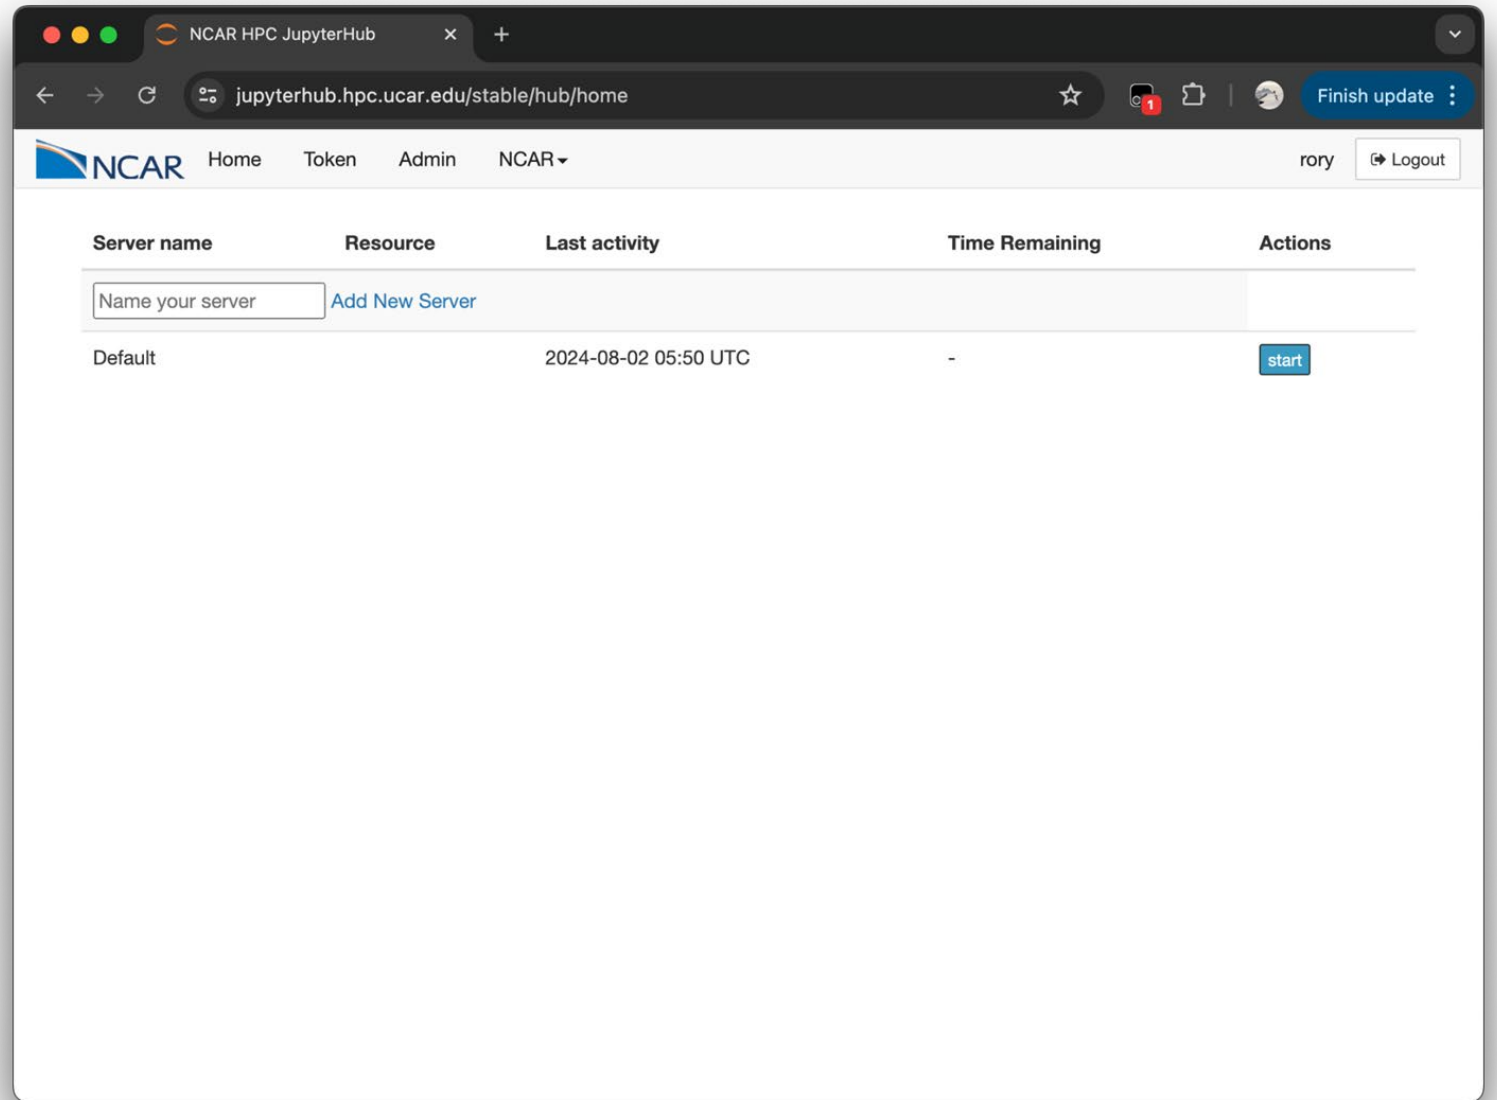
Jupyterhub - Logging In

Log in with your standard Duo credentials and respond to the push notification



Jupyterhub - Logging In

Start a server if you do not already have one running

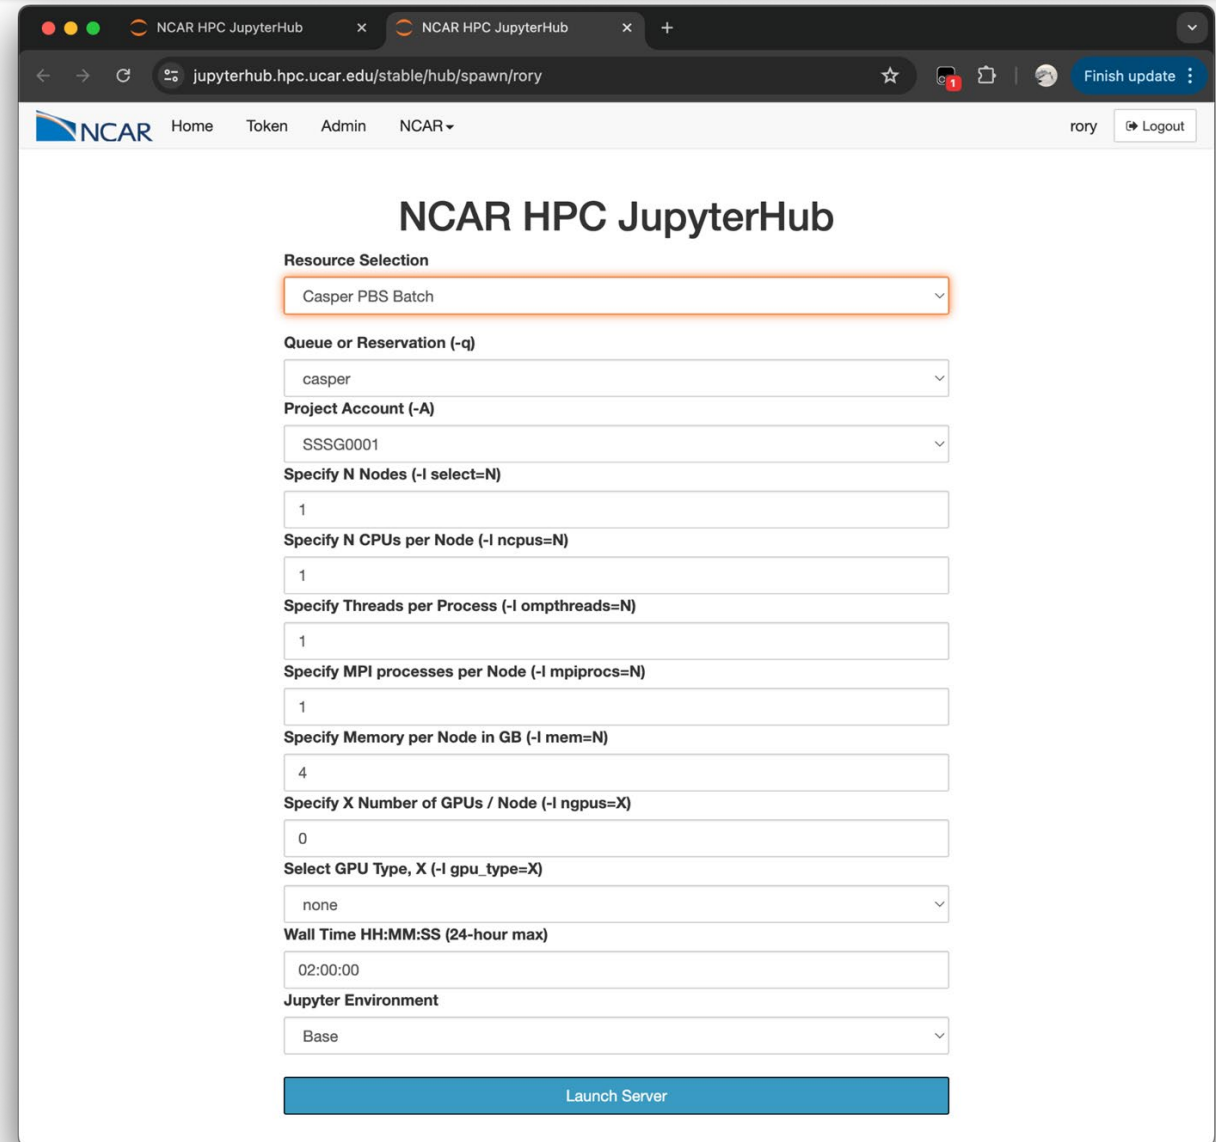


The screenshot shows the NCAR HPC JupyterHub interface in a web browser. The browser tab is titled "NCAR HPC JupyterHub" and the address bar shows "jupyterhub.hpc.ucar.edu/stable/hub/home". The page header includes the NCAR logo, navigation links for "Home", "Token", "Admin", and "NCAR", and a user profile for "rory" with a "Logout" button. Below the header is a table with the following columns: "Server name", "Resource", "Last activity", "Time Remaining", and "Actions".

Server name	Resource	Last activity	Time Remaining	Actions
<input type="text" value="Name your server"/>	Add New Server			
Default		2024-08-02 05:50 UTC	-	start

Jupyterhub - Logging In

Select a Resource and fill in batch request details, including number of cores and project account

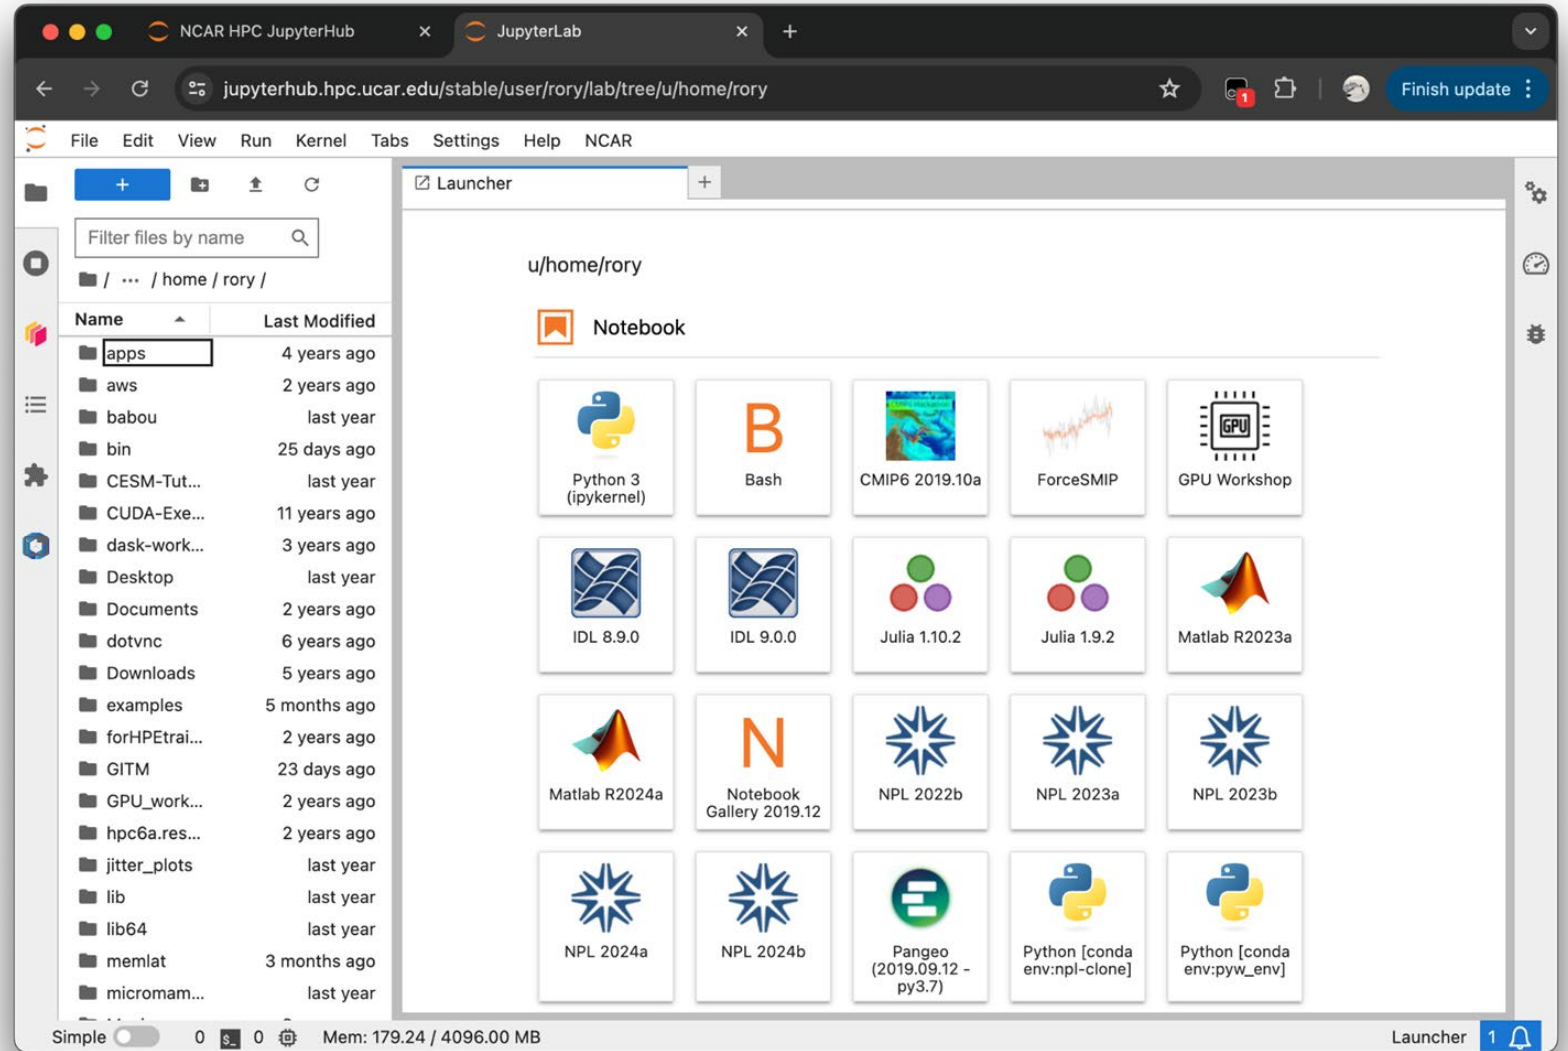


The screenshot shows the NCAR HPC JupyterHub interface. The browser address bar displays `jupyterhub.hpc.ucar.edu/stable/hub/spawn/rory`. The page title is "NCAR HPC JupyterHub". The navigation menu includes "Home", "Token", "Admin", and "NCAR". The user "rory" is logged in, with a "Logout" button. The main content area is titled "NCAR HPC JupyterHub" and contains a "Resource Selection" section with a dropdown menu set to "Casper PBS Batch". Below this are several input fields for request details: "Queue or Reservation (-q)" set to "casper", "Project Account (-A)" set to "SSSG0001", "Specify N Nodes (-l select=N)" set to "1", "Specify N CPUs per Node (-l ncpus=N)" set to "1", "Specify Threads per Process (-l ompthreads=N)" set to "1", "Specify MPI processes per Node (-l mpiprocs=N)" set to "1", "Specify Memory per Node in GB (-l mem=N)" set to "4", "Specify X Number of GPUs / Node (-l ngpus=X)" set to "0", "Select GPU Type, X (-l gpu_type=X)" set to "none", and "Wall Time HH:MM:SS (24-hour max)" set to "02:00:00". The "Jupyter Environment" dropdown is set to "Base". A blue "Launch Server" button is at the bottom.



Jupyterhub - Logging In

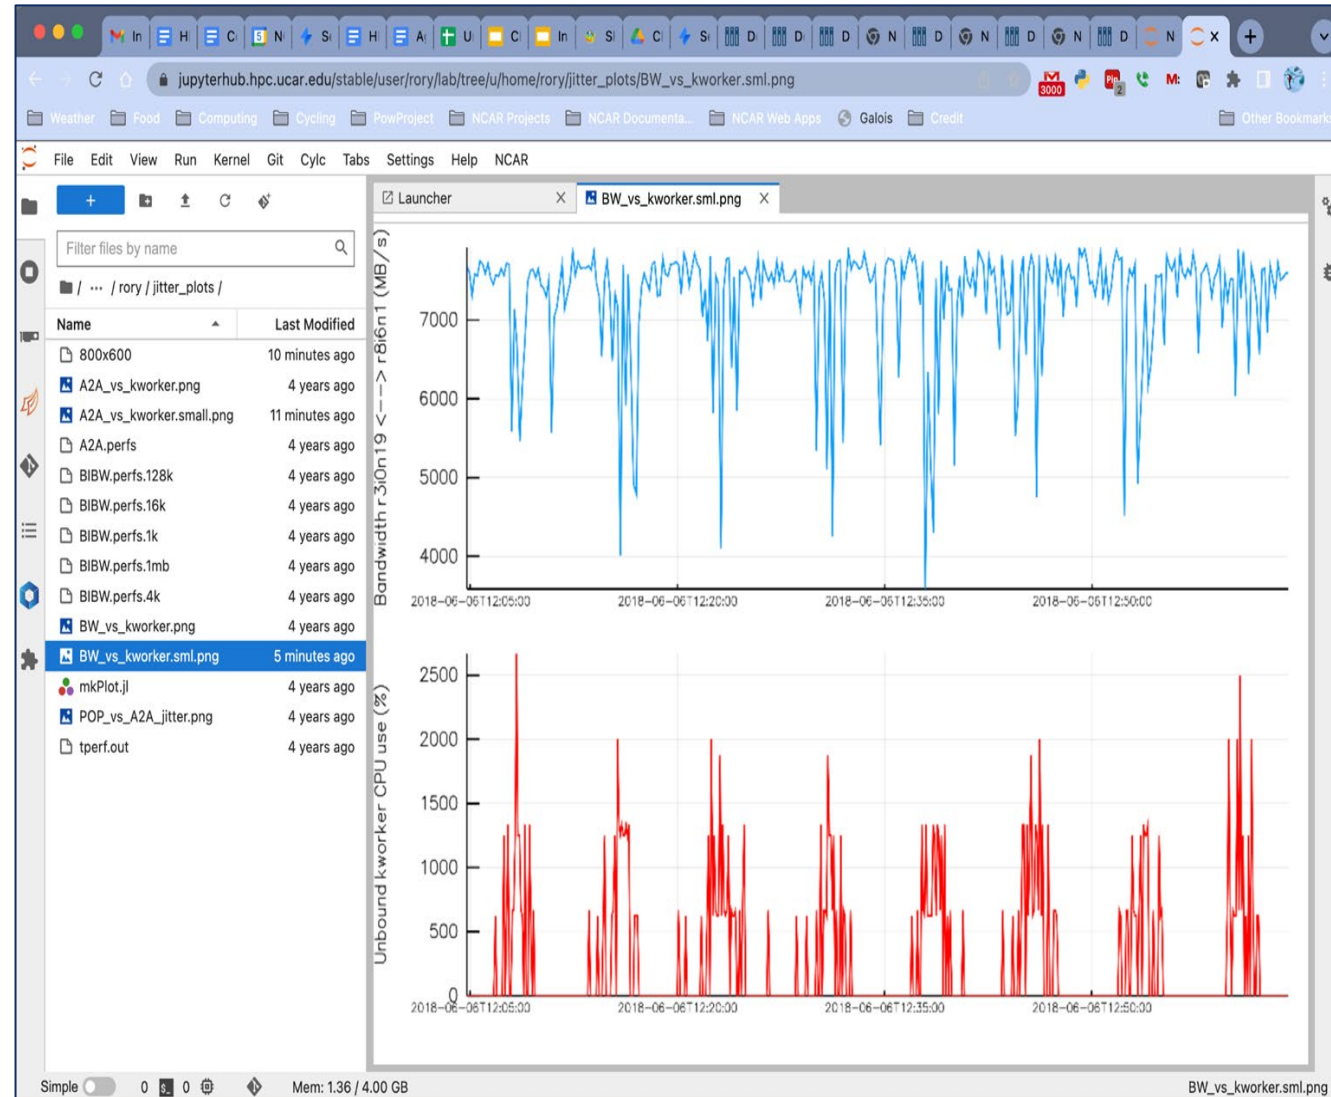
Many Notebooks and applications can be run from the main jupyterhub launcher screen.



Simple GUI Tasks with Jupyterhub

Jupyterhub is used to create sophisticated interactive computational notebooks for analysis, education, etc.

It can also be used for really unsophisticated things like viewing images



Be a Good Neighbor on Shared Resources

- Your activities coexists with those of other users
- CPUs and memory are shared on the login nodes
- Limit your usage to:
 - Reading and editing text/code
 - Compiling small programs
 - Performing data transfers
 - Interacting with the job scheduler
- Programs that use excessive resources on the login nodes will be limited (you will receive automated warnings via email if you exceed usage thresholds)
- In extreme cases, your processes may be terminated

Data Storage - GLADE

File spaces accessible from all HPC systems

File space	Quota	Backup	Uses
Home /glade/u/home/\$USER	50 GB	Yes	Settings, code, scripts
Work /glade/work/\$USER	2 TB	No	Compiled codes, models
Scratch /glade/derecho/scratch/\$USER	30 TB	Purged!	Run directories, temp output
Campaign /glade/campaign	N/A	No	Project space allocations

Keep track of usage with “gladequota”



Moving data to and from GLADE

- For short transfers use **scp/sftp** to transfer files
- For large transfers use **Globus**
 - To use Globus, create a Globus ID if you need an account, and search for **NCAR GLADE** or **NCAR Campaign Storage** endpoints
 - CISL endpoints currently can be activated for up to 30 -days
 - Globus has a web interface and a command -line interface
 - **Globus Connect Personal** can manage transfers from your local workstation as well



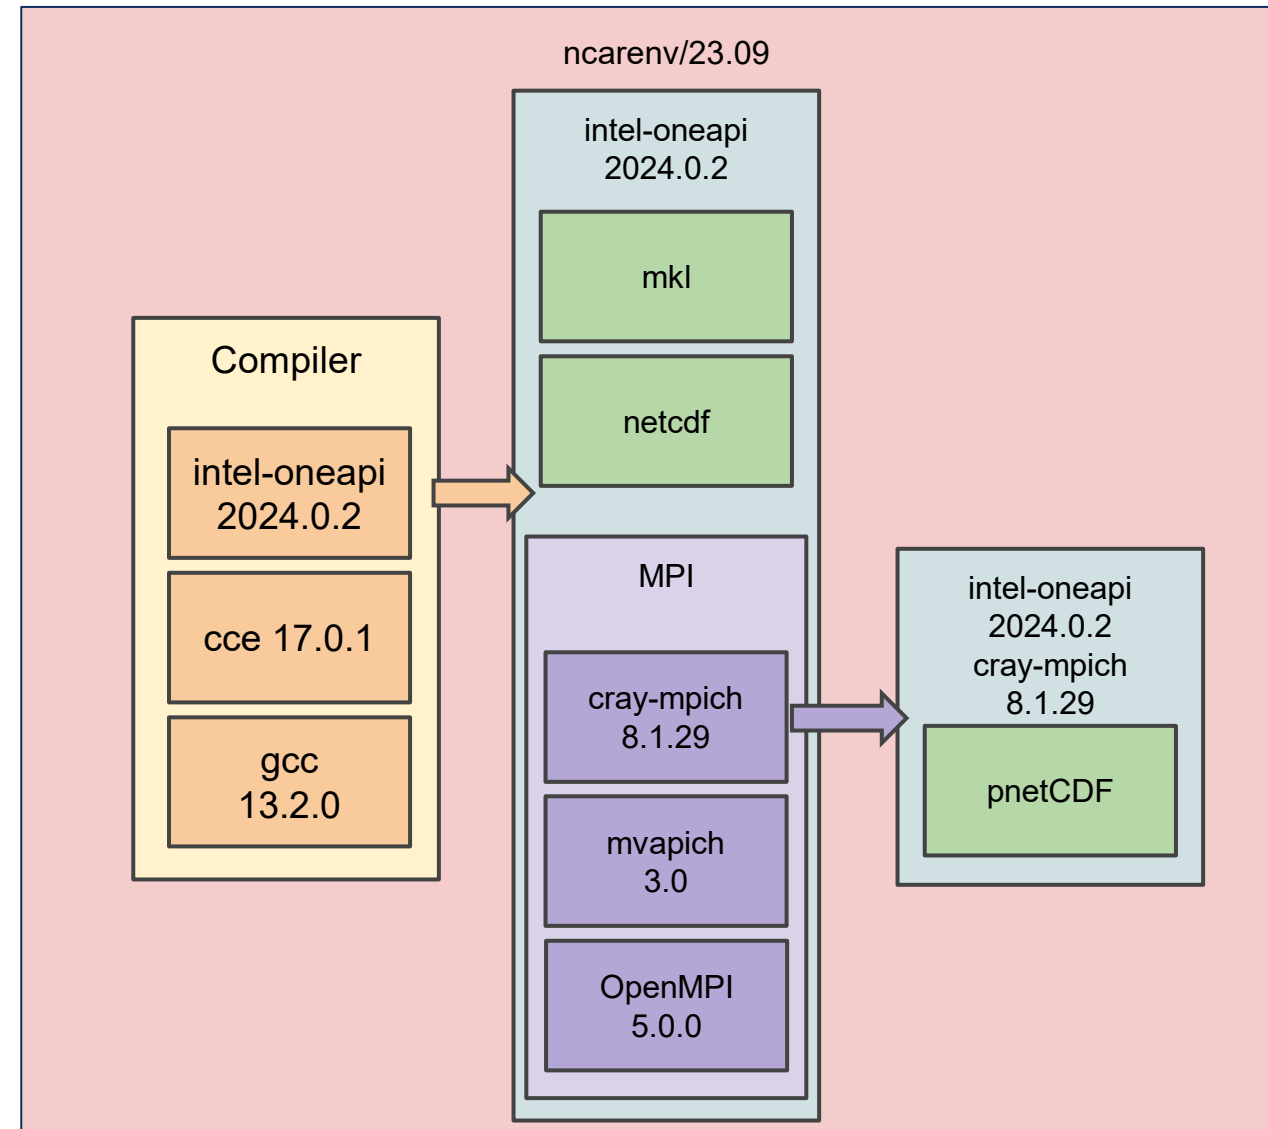
Environment Modules

- CISL installed software is provided as modules
- Modules provide access to runnable applications (compilers, debuggers, ...) as well as libraries (NetCDF, MPI, ...)
- Modules help ensure that all the software in your environment is mutually compatible
- **Note that Derecho and Casper each have independent collections of modules due to differences in their base software stacks. We make an effort to keep the stacks similar, but it is not possible in all cases.**



Using Modules

- **module load/unload <software>**
- **module avail**
 - show all currently-loadable modules
- **module list**
 - show loaded modules
- **module (--force) purge**
 - remove all loaded modules
- **module save/restore <name>**
 - create/load a saved software set
- **module spider <software>**
 - search for a particular module



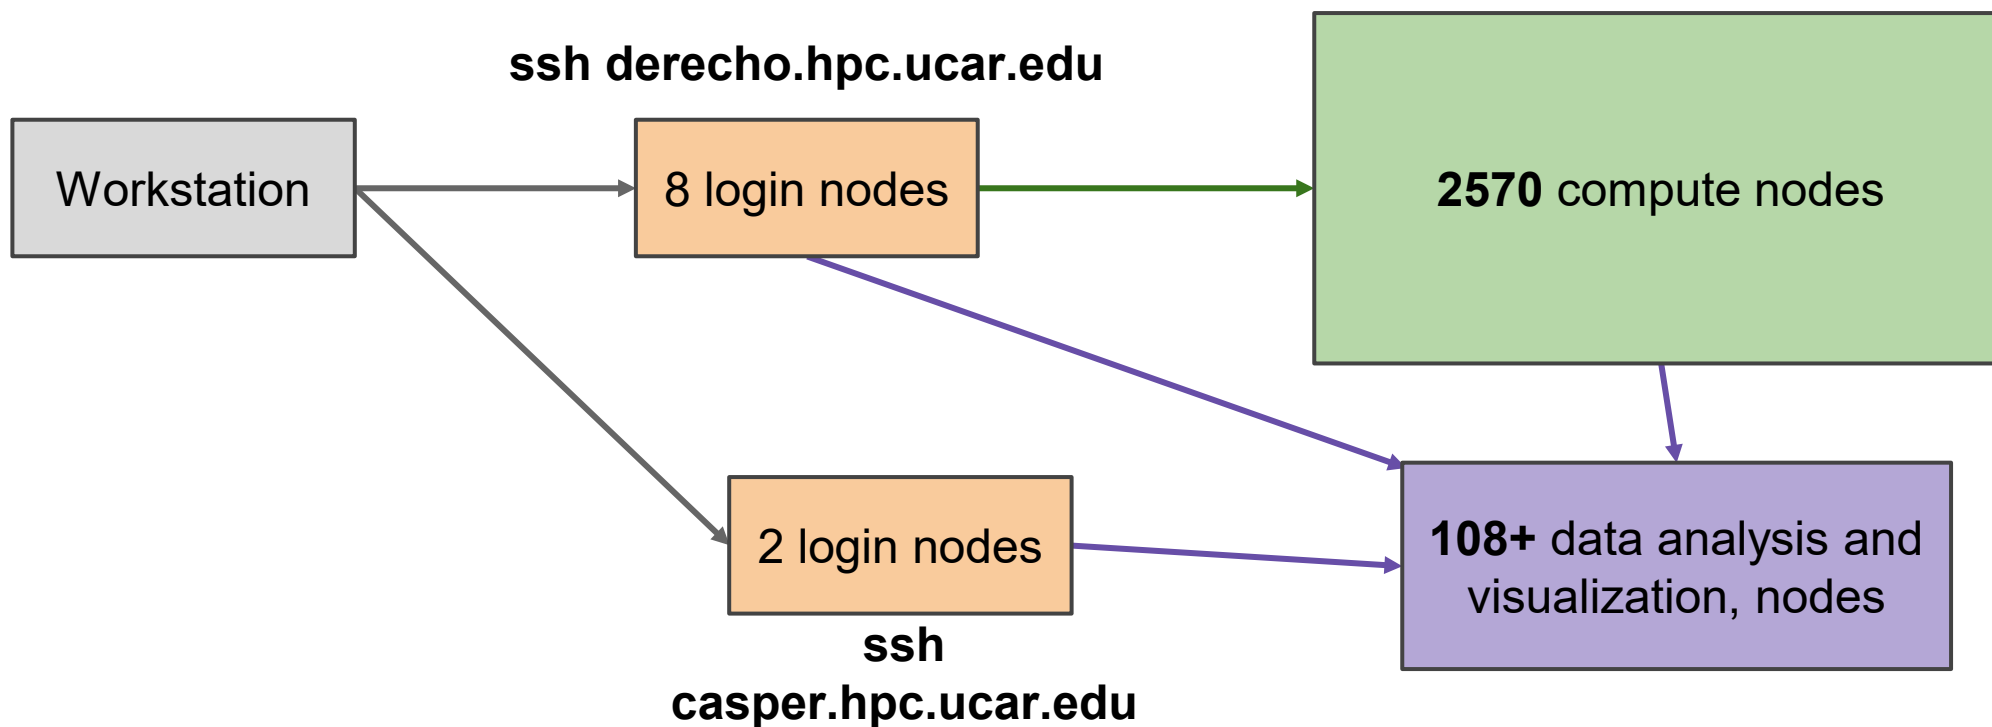
Available Software

- Compilers (Intel, CCE, GCC, NVHPC)
- Debuggers / Performance Tools (Linaro Forge, DDT, MAP, CrayPAT)
- MPI Libraries (Cray-MPICH, MVAPICH*, OpenMPI*, IntelMPI*)
- IO Libraries (NetCDF, PNetCDF, HDF5, Parallel IO)
- Analysis Languages (Python, Julia, R, IDL, Matlab)
- Convenience Tools (ncarccompilers, parallel, rclone)
- Many more ...



Using Batch Jobs and Compute Nodes

- Most tasks require too many resources to run on a login node
- Schedule these tasks to run on compute nodes using the **PBS** batch system



Derecho and Casper use separate allocations!

Using Batch Jobs and Compute Nodes

- Jobs request a given amount of compute resources, for an estimated amount of time, on specific hardware
- Jobs use core-hours, which are charged against your selected project/account
 - Remaining resources are viewable in SAM (sam.ucar.edu)
- Temporary files are often written by programs - set TMPDIR variable to scratch space to avoid job failures

Example PBS Batch Script

```
$ cat basic_mpi.pbs
#!/bin/bash
#PBS -N hello_pbs
#PBS -A <project_code>
#PBS -j oe
#PBS -k eod
#PBS -q main
#PBS -l walltime=00:05:00
#PBS -l select=8:ncpus=128:mpiprocs=128

### Set temp to scratch
export TMPDIR=/glade/derecho/scratch/${USER}/tmp
mkdir -p $TMPDIR

module --force purge
module load ncarenv/23.09 intel-oneapi/2024.0.2 ncarcompilers/1.0.0 cray-mpich/8.1.27

### Run MPT MPI Program
mpibind ./hello_world
```



Interacting with the PBS Scheduler

qsub <script> - submit batch job

qstat <jobid> - query job status

qdel <jobid> - delete/kill a job

qinteractive -A <project>

Run an interactive job

qcmd -A <project> -- <command>

Run command on a compute node

Customize Your Environment with Default Modules

- If you commonly load certain modules, you may wish to have them load automatically when logging onto a cluster
- The right way to do so is with saved module sets:
 `module --force purge`
 `module load ncarenv/23.09 intel -oneapi/2024.0.2`
 `module load ncarcompilers/1.0.0 cray -mpich/8.1.27`
 `module load conda/latest`
 `module save default`
- Make multiple sets and load them using `module restore <set>`
- **Don't put module load commands in your shell startup files!**

Getting Help

HPC Documentation - <https://ncar-hpc-docs.readthedocs.io>

The screenshot shows the NCAR HPC Documentation website. The header includes the NCAR CISL logo, the title 'NCAR HPC Documentation', a search bar, and a GitHub repository link. The main content area features the NCAR Computational & Information Systems Lab logo and the title 'User Documentation for NSF NCAR High Performance Computing'. A paragraph describes the site's purpose. A 'Selected Links' section lists several articles, with red arrows pointing to 'Getting Help', 'Submit a ticket', and 'Password Help'. A footer section provides information on how to submit a help request, including a phone number.

NCAR CISL NCAR HPC Documentation

COMPUTATIONAL & INFORMATION SYSTEMS LAB

User Documentation for NSF NCAR High Performance Computing

This is the home of the user documentation for the NSF NCAR high-performance computing (HPC) and storage resources managed by CISL. It includes searchable information specific to HPC resources, storage systems, authentication procedures and others, as well as additional how-to articles and troubleshooting articles.

Selected Links

- [Getting Started](#)
- [Using Derecho](#)
- [Using Casper](#)
- [Using JupyterHub](#)
- [Getting Help](#)

Don't find what you need? Log in here to submit a help request: [NCAR Research Computing](#).

You need a CIT password to submit a request. Call **303-497-2400** if you don't have one.

Consultant on Duty

Submit a ticket

Password Help

