Introduction to NCAR HPC Systems

For CISL's 2023 SIParCS Cohort and all new NCAR HPC users

Mea Trahan
CISL Consulting Services Group
mtrahan@ucar.edu

https://arc.ucar.edu/knowledge_base_documentation

May 24, 2023

NCAR
NATIONAL CENTER FOR ATMOSPHERIC RESEARCH

This material is based upon work supported by the National Center for Atmospheric Research, which is a major facility sponsored by the National Science Foundation under Cooperative Agreement No. 1852977.

Participant Code of Conduct

Our Pledge

UCAR and NCAR are committed to providing a safe, productive, and welcoming environment for all participants in any conference, workshop, field project or project hosted or managed by UCAR, no matter what role they play or their background. This includes respectful treatment of everyone regardless of gender, gender identity or expression, sexual orientation, disability, physical appearance, age, body size, race, religion, national origin, ethnicity, level of experience, political affiliation, veteran status, pregnancy, genetic information, as well as any other characteristic protected under state or federal law. (*link*)

Expected Behaviors

- All participants are treated with respect and consideration, valuing a diversity of views and opinions
- Be considerate, respectful, and collaborative
- Communicate openly with respect, critiquing ideas rather than individuals and gracefully accepting criticism
- Acknowledging the contributions of others
- Avoid personal attacks directed toward other participants
- Be mindful of your surroundings and of your fellow participants
- Alert UCAR staff and suppliers/vendors if you notice a dangerous situation or someone in distress
- Respect the rules and policies of the project and venue

Thank you for joining us today.

Here are a few things to note before we really get started:

- This tutorial is being recorded and will be available on the CISL website within the next few days.
- If you have questions, please enter them in the chat.
- Please keep your computer audio or phone muted!
- Please turn off your Zoom video (to save bandwidth).

Outline

- HPC Systems
- Systems Accounting Manager
- System Access
- Data Storage Spaces
- Software Environment
- Batch Job Submission
- Data Analysis Resources
- Additional Resources

HPC Systems


NCAR-Wyoming
Supercomputing Center
Operated under the sponsorship of
the National Science Foundation
NSF
8120 Veta Drive

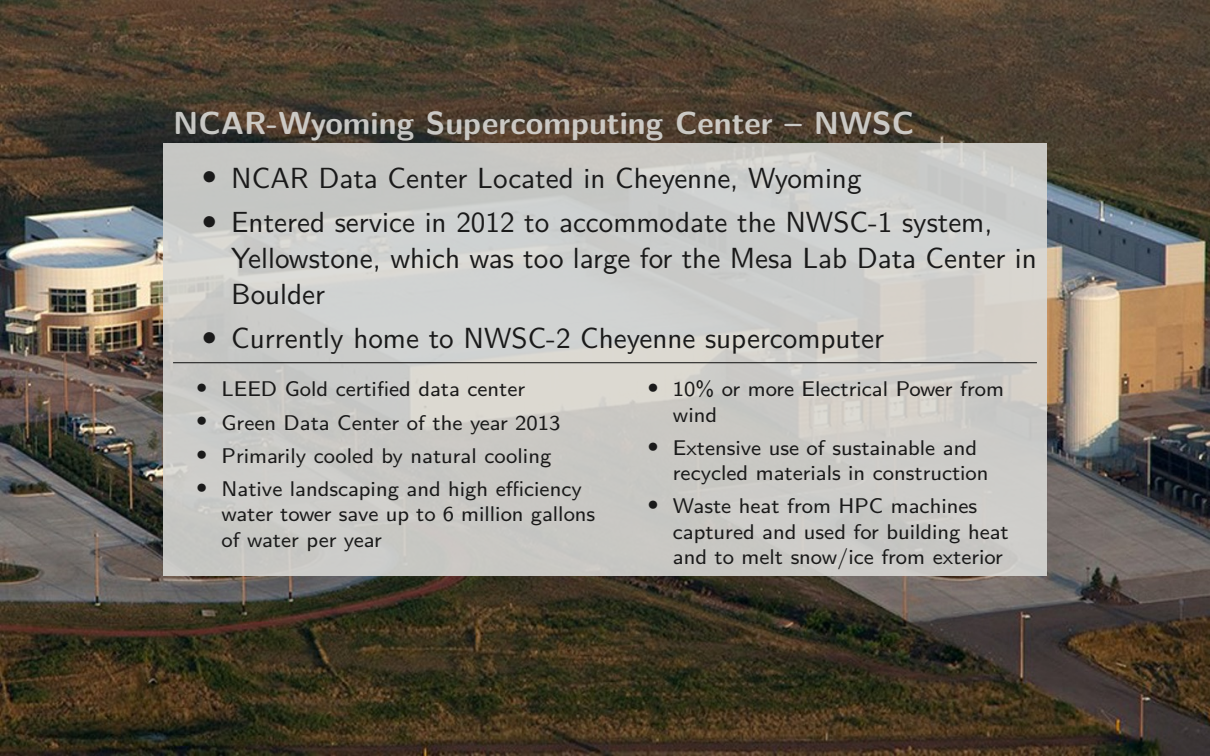

NCAR-Wyoming Supercomputing Center – NWSC

- NCAR Data Center Located in Cheyenne, Wyoming
- Entered service in 2012 to accommodate the NWSC-1 system, Yellowstone, which was too large for the Mesa Lab Data Center in Boulder
- Currently home to NWSC-2 Cheyenne supercomputer

- LEED Gold certified data center
- Green Data Center of the year 2013
- Primarily cooled by natural cooling
- Native landscaping and high efficiency water tower save up to 6 million gallons of water per year
- 10% or more Electrical Power from wind
- Extensive use of sustainable and recycled materials in construction
- Waste heat from HPC machines captured and used for building heat and to melt snow/ice from exterior


CHEYENNE

Cheyenne

SGI ICE XA Supercomputer

- 2nd supercomputing system deployed at NWSC
- Entered production January 2017
- Debuted at #21 of the World's Top 500 supercomputers, #100 in November '21

4032 Compute nodes (145,152 total cores)

- Dual socket, 18 cores per socket
- 2.3 GHz Intel Xeon (Broadwell) processors
- 313 TB total system memory, DDR4-2400
  - 64 GB/node, single-rank DIMM, 3,168 nodes
  - 128 GB/node, dual-rank DIMM, 864 nodes
- Mellanox EDR InfiniBand, Partial 9D Enhanced Hypercube Topology

6 login nodes

- Dual socket, 18 cores per socket, 256 GB memory/node

Documentation Home Page:

`https://arc.ucar.edu/knowledge_base/70549542`


Casper

WELCOME
TO
Casper
WYOCITY

Casper

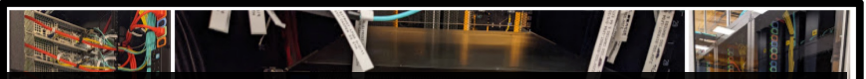

100 heterogeneous compute nodes of specialized nodes targeting data analysis, visualization, and GPU computing.

- 75 high-throughput computing (HTC) nodes across 2 generations of hardware for small computing tasks using 1 or 2 CPUs.
  - Typical HTC nodes have 384 GB system memory.
- 2 large memory nodes have 1.5 TB system memory.
- 9 nodes for data analysis and visualization jobs.
  - These nodes include a single NVIDIA Quadro GP100 16 GB GPU and 384 GB system memory.
- 10 *(+6 additional incoming)* nodes feature large-memory, dense GPU configurations to support explorations in machine learning (ML) and general purpose GPU computing.
  - 4 of these nodes feature 4 NVIDIA Tesla V100 32 GB GPUs and 768 GB system memory.
  - 6 of these nodes feature 8 NVIDIA Tesla V100 32 GB GPUs and 1 TB system memory.
  - *6 of these nodes **will** feature 4 NVIDIA Tesla A100 80 GB GPUs and 1 TB system memory.*
- 4 nodes are reserved for Research Data Archive workflows.

Documentation Home Page:

`https://arc.ucar.edu/knowledge_base/70549550`


DERECHO

Derecho

HPE/Cray EX Supercomputer

- 3rd supercomputing system deployed at NWSC
- Production planned for August 2023

2488 CPU + 82 GPU Compute nodes (323,712 total CPU cores)

- AMD EPYC™ 7763 Milan processors
- Cray Slingshot 11 Dragonfly Network Topology
- CPU Nodes:
  - Dual socket, 64 cores per socket
  - 256 GB DDR4 system memory
  - Single 200 Gb/sec Cassini network interface card
- GPU Nodes:
  - Single socket, 64 cores
  - 4× NVIDIA 1.41 GHz A100 Tensor Core GPUs per node and 600 GB/sec NVIDIA NVLink GPU intercont
  - 512 GB DDR4 system memory
  - 4× 200 Gb/sec Cassini network interface cards

8 login nodes

- 6 CPU-only, 2 GPU nodes with 2× NVIDIA A100s each

Documentation Home Page:

`https://arc.ucar.edu/knowledge_base/74317833`

SAM – Systems Accounting Manager

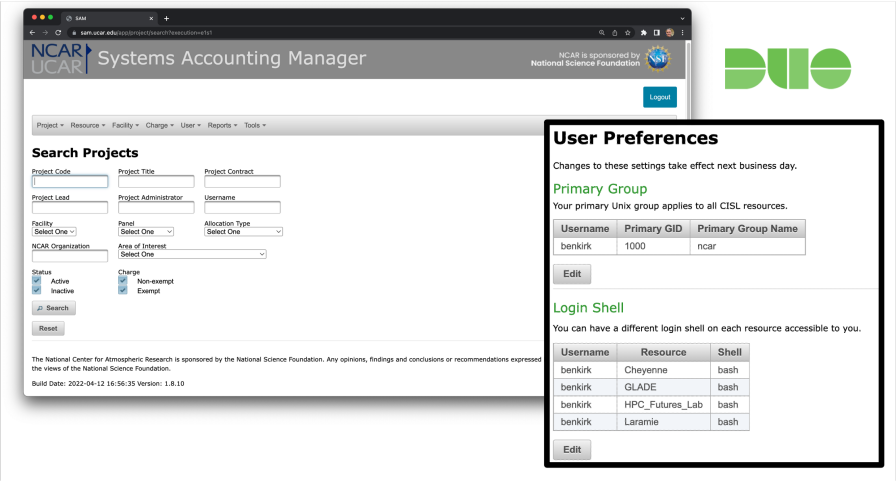

- Duo login at `https://sam.ucar.edu`
- Change some user settings (default project, shell, etc…)
- Query information about available projects and remaining allocation balance
- See history of jobs and charges

Logging in to Cheyenne or Casper

- Use `ssh` along with your username to log in

```
ssh -Y username@cheyenne.ucar.edu
ssh -Y username@casper.ucar.edu
```

- Use password+Duo for two-factor authentication
- You will be placed on a login node
  - Cheyenne - 6 login nodes: `cheyenne[1-6]`
  - Casper - 2 login nodes: `casper-login[1,2]`

More Details: `https://arc.ucar.edu/knowledge_base/74317885`

HPC System Access: ssh

```
$ ssh -Y benkirk@casper.ucar.edu
TokenResponse: # <---- Password here, then Duo Push Authentication
```

HPC System Access: ssh

```
$ ssh -Y benkirk@casper.ucar.edu
TokenResponse: # <---- Password here, then Duo Push Authentication
```

- **`ssh -Yvv username@cheyenne.ucar.edu`**
  for additional debugging output
- For additional 2-factor authentication options see
  `https://arc.ucar.edu/knowledge_base/70549637`

HPC System Access: ssh

```
$ ssh -Y benkirk@casper.ucar.edu
TokenResponse: # <---- Password here, then Duo Push Authentication
Last login: Tue Apr 19 07:39:20 2022 from 47.160.172.205

*******************************************************************************
*                   Welcome to Casper - April 18, 2022
*******************************************************************************
            Today in the Daily Bulletin (arc.ucar.edu)

     - No system downtimes for week of April 18-22
     - Default module updates completed April 11
     - Reminder: HPC systems and NWSC maintenance downtime in May

Documentation:         https://bit.ly/CISL-user-documentation
Key module commands:   module list, module avail, module spider, module help
CISL Help:             support.ucar.edu -- 303-497-2400
-------------------------------------------------------------------------------

benkirk@casper-login2(1)$
```

Good Citizenship

- **Be mindful of your usage on shared resources like the login nodes**
- Your activities coexists with those of other users
- CPUs and memory are shared on the login nodes
- Limit your usage to:
  - Reading and writing text/code
  - Compiling smaller programs
  - Performing data transfers
  - Interacting with the job scheduler
- Programs that use excessive resources on the login nodes will be terminated
- **Please do not attempt to run `sudo` on any CISL managed systems**
  - If you need help with a system issue, software installation request, etc. please ask for help!
  - Research Computing Help Desk: `https://rchelp.ucar.edu`

GLADE: GLobally Accessible Data Environment

- File spaces optimized for parallel IO, 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`	1 TB	No	Compiled codes, models
Scratch `/glade/scratch/$USER`	10 TB	**Purged!**	Run directories, temp output. **Purged at 120 days**
Project `/glade/p/entity/project_code`	N/A	No	Project space allocations

- `$HOME` is the only user file space that is regularly backed up.
  - *Snapshots* are also available for self directed file restoration. See `snapls`
- Check usage vs. quota with `gladequota`

GLADE: $HOME: snapls

```
# list $HOME directory snapshots
$ cd $HOME && snapls | grep snapshots | sort -r
 ... Apr 20 10:43 /glade/u/home/.snapshots/20220420-130001/benkirk
 ... Apr 20 10:43 /glade/u/home/.snapshots/20220420-120001/benkirk
 ... Apr 20 10:43 /glade/u/home/.snapshots/20220420-110001/benkirk
 ... Apr 20 08:35 /glade/u/home/.snapshots/20220420-090002/benkirk
 ... Apr 19 15:09 /glade/u/home/.snapshots/20220420-060001/benkirk
 ... Apr 19 15:09 /glade/u/home/.snapshots/20220420-030001/benkirk
 ... Apr 19 15:09 /glade/u/home/.snapshots/20220420-000001/benkirk
 ... Apr 19 15:09 /glade/u/home/.snapshots/20220419-180002/benkirk
 ... Apr 19 11:08 /glade/u/home/.snapshots/20220419-120001/benkirk
 ... Apr 18 08:34 /glade/u/home/.snapshots/20220419-000001/benkirk
 ... Apr 14 15:12 /glade/u/home/.snapshots/20220418-000001/benkirk
 ... Apr 14 15:12 /glade/u/home/.snapshots/20220417-000001/benkirk
 ... Apr 14 15:12 /glade/u/home/.snapshots/20220416-000001/benkirk
 ... Apr 14 15:12 /glade/u/home/.snapshots/20220415-000001/benkirk
 ... Apr 13 19:38 /glade/u/home/.snapshots/20220414-000001/benkirk
 # peruse a snapshot ...
 $ ls /glade/u/home/.snapshots/20220420-000001/benkirk
```

GLADE: $HOME: snapls

```
# list $HOME directory snapshots
$ cd $HOME && snapls | grep snapshots | sort -r
 ... Apr 20 10:43 /glade/u/home/.snapshots/20220420-130001/benkirk
 ... Apr 20 10:43 /glade/u/home/.snapshots/20220420-120001/benkirk
 ... Apr 20 10:43 /glade/u/home/.snapshots/20220420-110001/benkirk
 ... Apr 20 08:35 /glade/u/home/.snapshots/20220420-090002/benkirk
 ... Apr 19 15:09 /glade/u/home/.snapshots/20220420-060001/benkirk
 ... Apr 19 15:09 /glade/u/home/.snapshots/20220420-030001/benkirk
 ... Apr 19 15:09 /glade/u/home/.snapshots/20220420-000001/benkirk
 ... Apr 19 15:09 /glade/u/home/.snapshots/20220419-180002/benkirk
 ... Apr 19 11:08 /glade/u/home/.snapshots/20220419-120001/benkirk
 ... Apr 18 08:34 /glade/u/home/.snapshots/20220419-000001/benkirk
 ... Apr 14 15:12 /glade/u/home/.snapshots/20220418-000001/benkirk
 ... Apr 14 15:12 /glade/u/home/.snapshots/20220417-000001/benkirk
 ... Apr 14 15:12 /glade/u/home/.snapshots/20220416-000001/benkirk
 ... Apr 14 15:12 /glade/u/home/.snapshots/20220415-000001/benkirk
 ... Apr 13 19:38 /glade/u/home/.snapshots/20220414-000001/benkirk
 # peruse a snapshot ...
 $ ls /glade/u/home/.snapshots/20220420-000001/benkirk
```

GLADE: gladequota

```
$ gladequota
Current GLADE space usage: benkirk

  Space                                      Used        Quota     % Full      # Files
------------------------------------------ ---------- ---------- --------- -----------
/glade/scratch/benkirk                     10.53 GiB   10.00 TiB    0.10 %      403423
/glade/work/benkirk                         7.30 GiB 1024.00 GiB    0.71 %      143469
/glade/u/home/benkirk                       0.77 GiB  100.00 GiB    0.77 %        5364
------------------------------------------ ---------- ---------- --------- -----------
/glade/u/sampledata                        51.82 GiB 1024.00 GiB    5.06 %         104
/glade/u/cesm-scripts                     252.86 GiB 1024.00 GiB   24.69 %      273236
/glade/p/cesm                            1155.98 TiB 1200.00 TiB   96.33 %    11456021
------------------------------------------ ---------- ---------- --------- -----------
Campaign: benkirk (user total)              0.00 GiB          n/a       n/a           6
/glade/campaign/collections/cmip/CMIP6   3800.64 TiB 4096.00 TiB   92.79 %     4376451
/glade/campaign/cesm                     8758.77 TiB 10240.00 TiB  85.53 %    30425283
/glade/campaign/cgd/cesm                 1874.32 TiB 2048.00 TiB   91.52 %     5692283
/glade/campaign/cisl/csg                  556.75 GiB   23.00 TiB    2.36 %        1750

/glade/scratch  - 78.7% used (12085 TiB used out of 15360 TiB total)
```

Note: $HOME quota is reported **2**× due to backup/redundancy implementation.

Campaign Storage

- Resource for storing data on publication timescales
- Multiple access methods:
  - Globus (NCAR Campaign Storage)
  - Casper nodes (`/glade/campaign/`)
  - Data access nodes (`/glade/campaign/`)
- Allocated to and managed by NCAR labs and can be requested by University users

More Details & Request Process: `https://arc.ucar.edu/knowledge_base/70549621`

Collections

- Curated data collections available on Cheyenne and Casper to facilitate easy access to research data sets
- **RDA**: Research Data Archive
  - `/glade/collections/rda/`

    `https://rda.ucar.edu/`
- **CMIP6**: Coupled Model Intercomparison Project
  - `/glade/collections/cmip/CMIP6/`

    `https://www2.cisl.ucar.edu/computing-data/data/cmip-analysis-platform`

Data Transfer

For *short, small* transfers use **scp/sftp** or **rsync** to transfer files

For *long, 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

Available Software

CISL provides a wide range of software tools for use inside the HPC environment:

- Compilers (Intel, GNU, PGI)
- Arm Forge Debuggers / Performance Tools (DDT, MAP)
- MPI Libraries (MPT, Intel MPI, OpenMPI)
- IO Libraries (NetCDF, PNetCDF, HDF5)
- Analysis Languages (Python, Julia, R, IDL, Matlab)
- Convenience Tools (ncarcompilers, parallel, rclone)
- Many more: `https://arc.ucar.edu/knowledge_base/70549892`

Need something else? Submit a request at `https://rchelp.ucar.edu`

Environment Modules

- CISL installed software is provided as modules
- Modules provide access to runnable applications (compilers, debuggers, ...) as well as libraries (NetCDF, MPI, ...)
- Modules prevent loading incompatible software into your environment

Note that Cheyenne and Casper each have independent collections of modules!

Module Commands

- **`module list`**
  Lists currently loaded modules
- **`module avail`**
  Shows all modules currently available (dynamic, depends on modules loaded)
- **`module load/unload <software>`**
  Loads or unloads the requested software package into the user environment
- **`module swap <software> <software/Other.Version>`**
  Switch to a different version of a software package
- **`module purge`**
  Removes all loaded modules
- **`module save/restore <name>`**
  Saves or loads a collection of modules
- **`module spider <software>`**
  Searches for particular software

More Details: `https://arc.ucar.edu/knowledge_base/72581272`

```
$ module avail
------------- /glade/u/apps/ch/modulefiles/default/compilers --------------
   ...
   gnu/9.1.0           intel/19.1.1 (L,D)      nvhpc/21.11
   gnu/10.1.0   (D)    intel/2022.1            nvhpc/22.1
   ...
   intel/18.0.5        nvhpc/21.3              pgi/20.4    (D)
---------------- /glade/u/apps/ch/modulefiles/default/idep -----------------
   ...
   arm-forge/21.1.1   (D)    matlab/R2021b
   arm-reports/19.1   (D)    matlab/R2022a       (D)
   arm-reports/20.0.2        nano/4.3
   ...
   ------------ /glade/u/apps/ch/modulefiles/default/intel/19.1.1 ------------
   ...
   esmf_libs/8.2.0 (D)    mpt/2.24
   fftw/3.3.8             mpt/2.25             (L,D)
   fftw/3.3.9      (D)    ncarcompilers/0.5.0 (L)
   gdal/3.0.4             ncl/6.6.2
   ...
   grib-api/1.28.0        openmpi/3.1.4
   grib-libs/1.2          openmpi/4.0.3
   ...
------- /glade/u/apps/ch/modulefiles/default/mpt/2.25/intel/19.1.1 --------
   fftw-mpi/3.3.9      pio/1.10.1        pio/2.5.6d    pio/2.5.7
   hdf5-mpi/1.10.8     pio/2.5.4  (D)    pio/2.5.6     pnetcdf/1.12.2
   netcdf-mpi/4.8.1    pio/2.5.5         pio/2.5.7d
```

Modules

```
$ module list

Currently Loaded Modules:
  1) ncarenv/1.3    3) ncarcompilers/0.5.0   5) netcdf/4.8.1
  2) intel/19.1.1   4) mpt/2.25

$ which icc && echo $NETCDF
/glade/u/apps/ch/opt/ncarcompilers/0.5.0/intel/19.1.1/icc
/glade/u/apps/ch/opt/netcdf/4.8.1/intel/19.1.1/
```

Modules

```
$ module swap intel intel/2022.1

Due to MODULEPATH changes, the following have been reloaded:
  1) mpt/2.25     2) ncarcompilers/0.5.0     3) netcdf/4.8.1

The following have been reloaded with a version change:
  1) intel/19.1.1 => intel/2022.1
```

Modules

```
$ module swap intel intel/2022.1

Due to MODULEPATH changes, the following have been reloaded:
  1) mpt/2.25     2) ncarcompilers/0.5.0     3) netcdf/4.8.1

The following have been reloaded with a version change:
  1) intel/19.1.1 => intel/2022.1
```

```
$ which icc && echo $NETCDF
/glade/u/apps/ch/opt/ncarcompilers/0.5.0/intel/2022.1/icc
/glade/u/apps/ch/opt/netcdf/4.8.1/intel/2022.1/

$ module list

Currently Loaded Modules:
  1) ncarenv/1.3    3) ncarcompilers/0.5.0   5) netcdf/4.8.1
  2) intel/2022.1   4) mpt/2.25
```

Changing Your Default Modules

Don't put module load commands in your shell startup files!

- If you commonly load certain modules, you may wish to have them load automatically when logging onto a cluster
- The proper way to do so is with saved module sets:

```
module load ncl python nco mkl
module save default
```

- You can make multiple named sets and load them using `module restore <set>`

Compiling Software

Considerations when compiling software

- Use **ncarcompilers** module along with library modules (e.g., `netcdf`) to simplify compiling and linking *(it adds include and link flags for you)*
- When using MPI, make sure you run with the same library with which you compiled your code
  - Strongly recommend loading desired modules inside run scripts, *more later*
- **We strongly recommend you build code on the machine on which you will run**
  - Cheyenne and Casper have different CPUs and operating systems

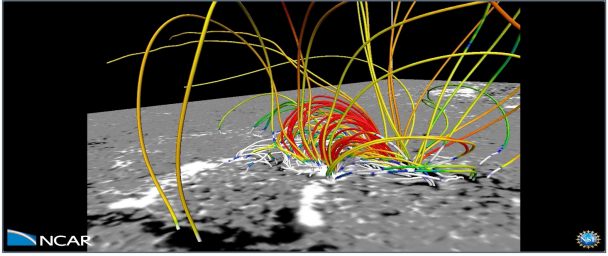

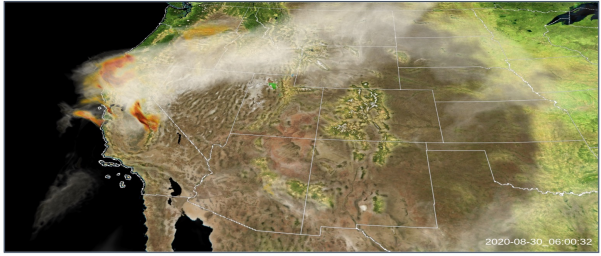

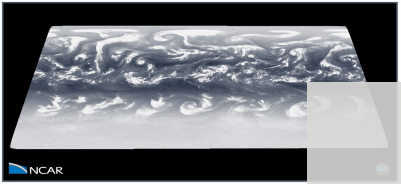

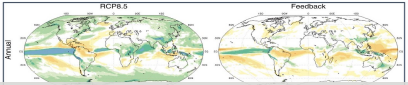

Running Jobs

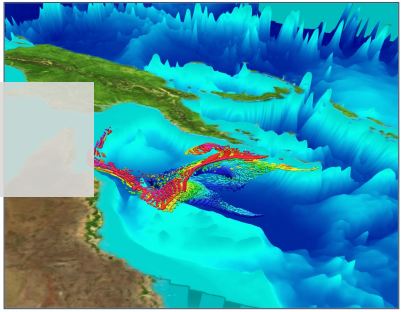

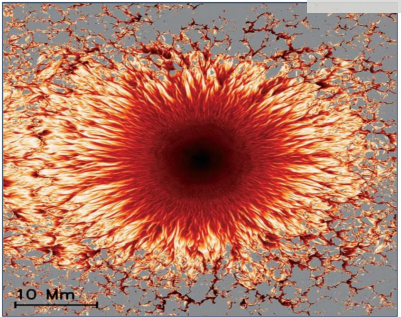

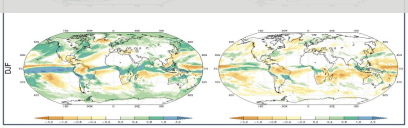

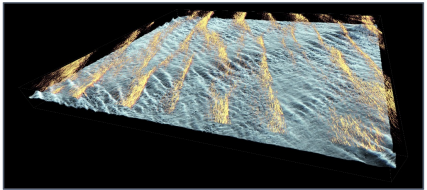

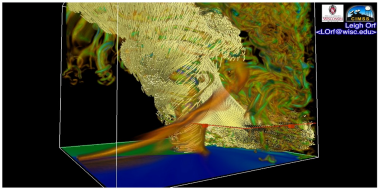

Accessing Compute nodes via Batch Jobs

Run large tasks on compute nodes using batch jobs

- Most tasks require too many resources to run on a login node
- Schedule these tasks to run on Cheyenne or Casper compute nodes using **Altair's PBS**
- Jobs request a given number of compute tasks for an estimated wall-time on specified hardware
- Jobs use core-hours, which are charged against your selected project/account
  - Remaining resources are viewable in SAM
- Temporary files are often written by programs - set TMPDIR variable to scratch space to avoid job failures

Altair | PBS Works™

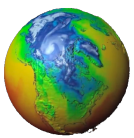

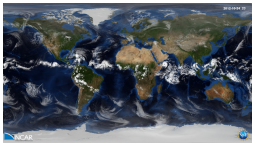

Example PBS batch job scripts

```
$ cat basic_mpi.pbs
#!/bin/bash
#PBS -N hello_pbs
#PBS -o pbsjob.log
#PBS -A <project_code>
#PBS -j oe
#PBS -k eod
#PBS -q regular
#PBS -l walltime=00:05:00
### Select 2 nodes with 36 CPUs each for a total of 72 MPI processes
#PBS -l select=2:ncpus=36:mpiprocs=36:ompthreads=1

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

module load mpt/2.25 && module list

### Interrogate Environment (optional, personal preference)
env | egrep "PBS|MPI|THREADS" | sort

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

More Examples:

- Cheyenne: https://arc.ucar.edu/knowledge_base/72581486
- Casper: https://arc.ucar.edu/knowledge_base/72581394

PBS Scheduler Interaction

- `qsub <script>`
  submit a batch job, see `man qsub`
- `qstat <jobid>`
  query job status
- `qinteractive -A <project> ...`
  run an interactive job with access to 1 or more CPUs
- `qcmd -A <project> -- cmd.exe`
  run a command on a single compute node
- `qhist`
  search PBS logs for finished jobs, see `qhist --help`

Example PBS batch job submission

```
# submit a batch script for execution with 'qsub'
$ qsub ./basic_pbs.sh
3864501.chadmin1.ib0.cheyenne.ucar.edu

# check the status of my running jobs with 'qstat' (may take up to 10s to appear)
$ qstat -u $USER
Req'd  Req'd   Elap
Job ID           Username Queue    Jobname    SessID NDS TSK Memory Time  S Time
---------------- -------- -------- ---------- ------ --- --- ------ ----- - -----
3864501.chadmi* benkirk  regular  hello_pbs      --    2  72    --  00:05 Q   --

# delete a job from the queue with 'qdel'
$ qdel 3864501

# check out my recently completed jobs with 'qhist'
$ qhist -u $USER
Job ID    User       Queue    Nodes NCPUs Finish   RMem(GB)  Mem(GB)   CPU(%)  Elap(h)
3865076   benkirk    regular      1    36 21-1434         -      1.0     92.8     0.20
3865065   benkirk    regular      1    36 21-1434         -      1.0     97.0     0.20
3865054   benkirk    shareex      1     1 21-1433         -      0.0      0.0     0.01
...

# submit a job from cheyenne to run on casper
cheyenne$ qsub -q casper@casper-pbs ./largemem_casper.sh
```

Example PBS batch job scripts

```
$ cat largemem_casper.sh
#PBS -N largemem_example
#PBS -o largemem-out.log
#PBS -A <project_code>
#PBS -j oe
#PBS -k eod
#PBS -q casper
#PBS -l walltime=00:02:00
### Select 24 OpenMP threads on 1 node with a total of 400GB RAM
#PBS -l select=1:ncpus=24:mpiprocs=1:ompthreads=24:mem=400G

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

module load ncl/6.6.2 && module list

### Interrogate Environment (optional, personal preference)
env | egrep "PBS|MPI|THREADS" | sort

### do something
echo "Hello from " $(hostname)
./hello_world && echo "Done."
```

More Examples:

- Cheyenne: https://arc.ucar.edu/knowledge_base/72581486
- Casper: https://arc.ucar.edu/knowledge_base/72581394

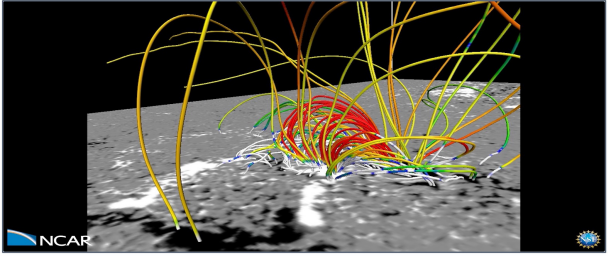

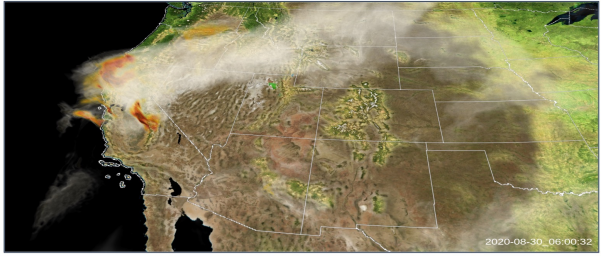

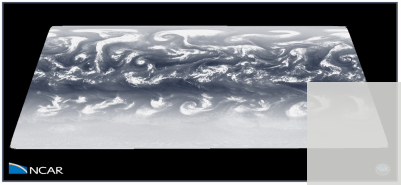

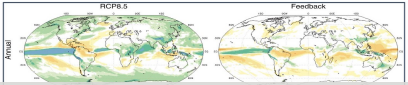

Data Analysis

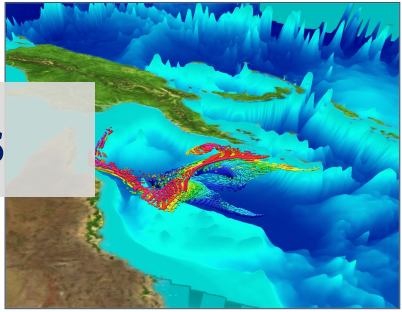

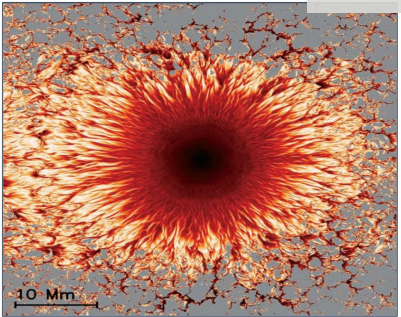

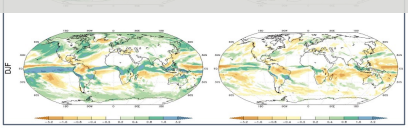

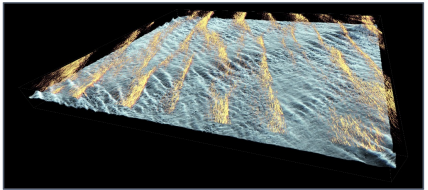

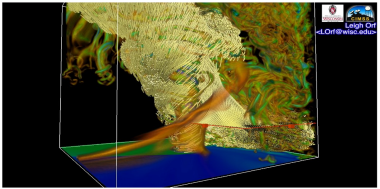

JupyterHub

- JupyterHub is a hosted Jupyter Notebook platform inside NCAR's HPC environment that is very useful for data analysis & processing workflows

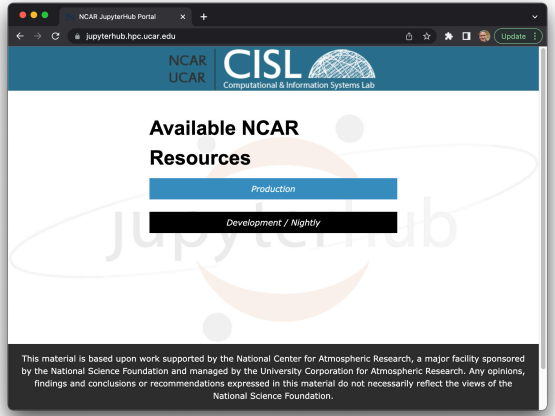

`https://jupyterhub.hpc.ucar.edu/`

More Details: `https://arc.ucar.edu/knowledge_base/70549913`

JupyterHub

- JupyterHub is a hosted Jupyter Notebook platform inside NCAR's HPC environment that is very useful for data analysis & processing workflows

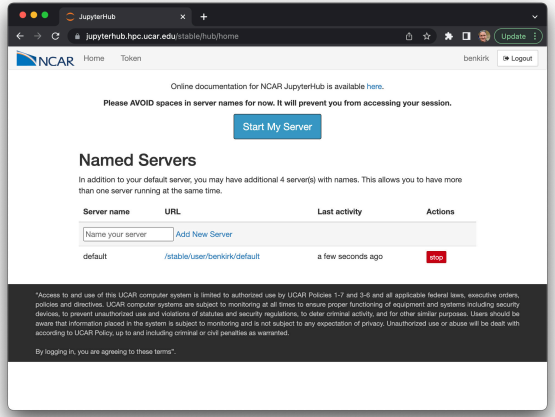

`https://jupyterhub.hpc.ucar.edu/`

More Details: `https://arc.ucar.edu/knowledge_base/70549913`

JupyterHub

- JupyterHub is a hosted Jupyter Notebook platform inside NCAR's HPC environment that is very useful for data analysis & processing workflows

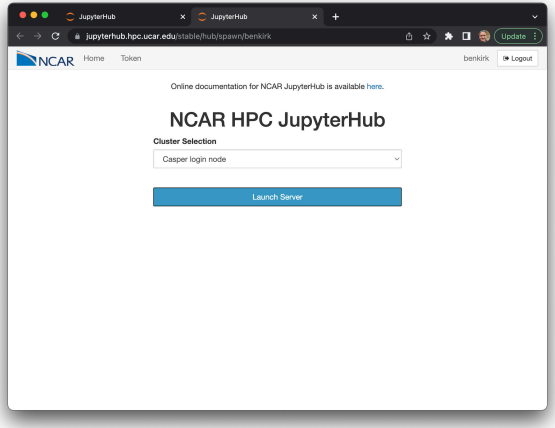

`https://jupyterhub.hpc.ucar.edu/`

More Details: `https://arc.ucar.edu/knowledge_base/70549913`

JupyterHub

- JupyterHub is a hosted Jupyter Notebook platform inside NCAR's HPC environment that is very useful for data analysis & processing workflows

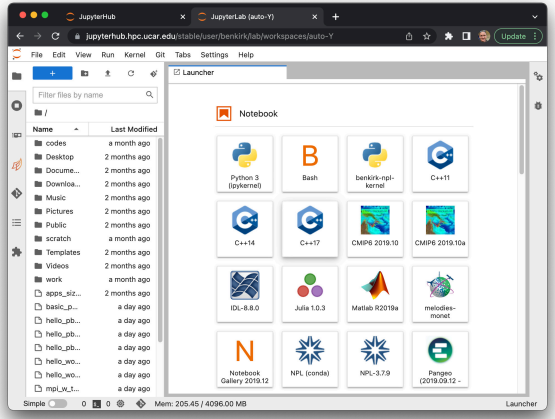

`https://jupyterhub.hpc.ucar.edu/`

More Details: `https://arc.ucar.edu/knowledge_base/70549913`

JupyterHub

- JupyterHub is a hosted Jupyter Notebook platform inside NCAR's HPC environment that is very useful for data analysis & processing workflows

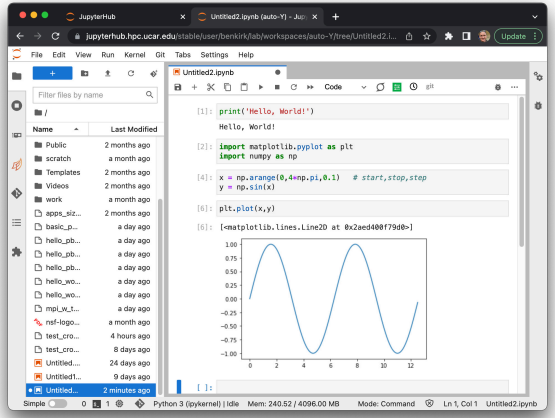

`https://jupyterhub.hpc.ucar.edu/`

More Details: `https://arc.ucar.edu/knowledge_base/70549913`

Remote Desktop & Graphical Programs: VNC

VNC can be used to run a remote KDE or GNOME desktop to support graphical applications.

```
# launch 'vncmngr' on casper, matched with a VNC client locally
casper$ vncmgr
```

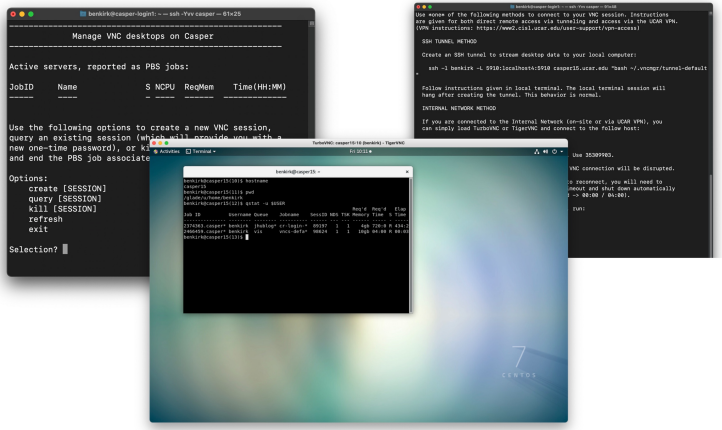

CISL recommends TigerVNC or TurboVNC client, see https://arc.ucar.edu/knowledge_base/72581380

Remote Desktop: FastX

- FastX is a alternate remote desktop service requiring only a web browser, or optional desktop client.
- **FastX is only accessible from the NCAR VPN**

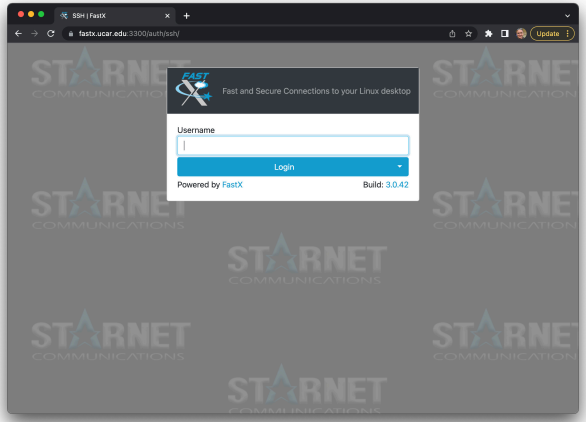

`https://fastx.ucar.edu:3300/`

More Details: `https://arc.ucar.edu/knowledge_base/72581391`

Remote Desktop: FastX

- FastX is a alternate remote desktop service requiring only a web browser, or optional desktop client.
- **FastX is only accessible from the NCAR VPN**

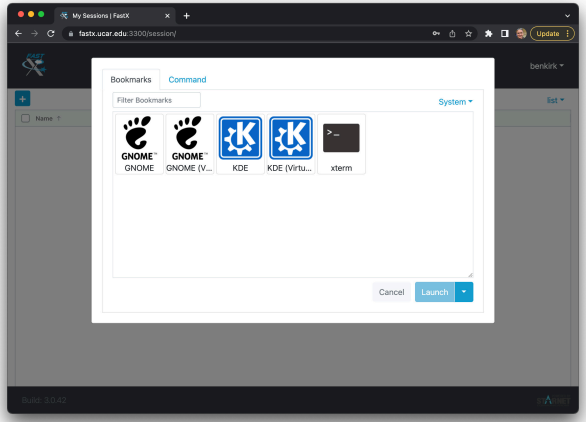

`https://fastx.ucar.edu:3300/`

More Details: `https://arc.ucar.edu/knowledge_base/72581391`

Remote Desktop: FastX

- FastX is a alternate remote desktop service requiring only a web browser, or optional desktop client.
- **FastX is only accessible from the NCAR VPN**

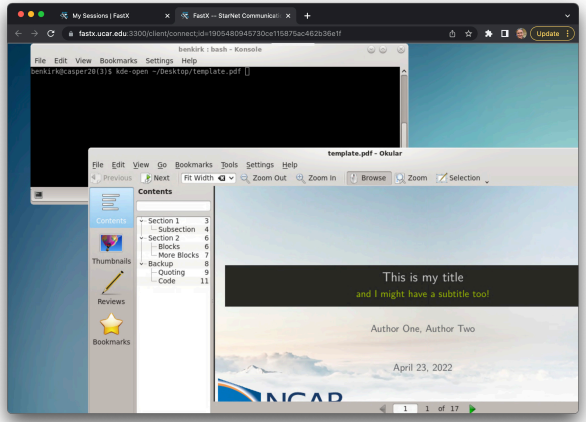

`https://fastx.ucar.edu:3300/`

More Details: `https://arc.ucar.edu/knowledge_base/72581391`

Remote Desktop: FastX

- FastX is a alternate remote desktop service requiring only a web browser, or optional desktop client.
- **FastX is only accessible from the NCAR VPN**

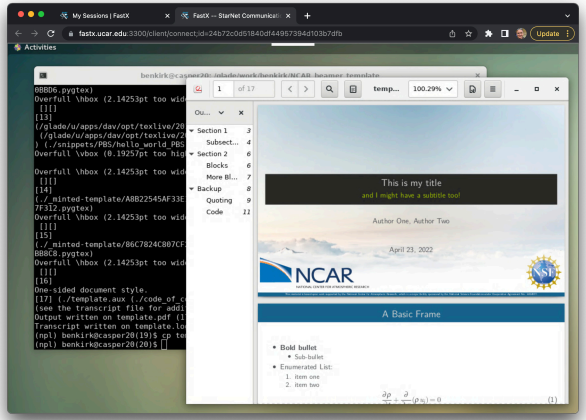

`https://fastx.ucar.edu:3300/`

More Details: `https://arc.ucar.edu/knowledge_base/72581391`

Going Further: Additional Resources & Requesting Help

- **Advanced Research Computing Documentation**:
  `https://arc.ucar.edu/knowledge_base_documentation`

- **CISL Help Desk**:
  `https://rchelp.ucar.edu`

  Submit a ticket to request help with a particular issue.

- **HPC Tutorials**:
  `https://www2.cisl.ucar.edu/what-we-do/training-library/hpc-tutorials`

  In-depth tutorials on numerous topics, including additional details on many of the items covered here today.
  - Introduction to NCAR HPC Systems
  - Job Scheduling with PBS Pro
  - JupyterHub at NCAR
  - NCAR Storage Spaces
  - Optimizing Resource Use in Scheduled Jobs
  - Remote desktop services on Casper
  - Starting Casper Jobs with PBS Pro
  - Using Globus at NCAR

Best Practices for Support Tickets

When submitting a support ticket please include as much detail as possible to enable quicker resolution:

- Resource name (Cheyenne, Casper, JupyterHub,...),
- **Exact** error messages and/or paths to error output,
- Batch script location,
- PBS JobID(s) of failed effort,
- Run & source directory paths (ideally UNIX-readable by 'others'),
- Any other pertinent information:
  - Last time this **exact** workflow was successful, if any (or changes since last success),
  - Troubleshooting steps already attempted, etc. ...
- ***And please remember to let us know when your issue is resolved!***

`https://rchelp.ucar.edu`

Questions?

Customizing your default environment

bash

```
$ cat  /.profile
alias rm"=rm -"i

# Add programs built for each cluster
if [[ $HOSTNAME == cheyenne* ]]; then
   export PATH=~/local/ch/bin:$PATH
else
   export PATH=~/local/dav/bin:$PATH
fi

# Settings for interactive shells
if [[ $- == *i* ]]; then
   PS1="\u@\h:\w> "
fi
```

tcsh

```
$ cat  /.tcshrc
alias rm "rm -"i

# Add programs built for each cluster
if ( $HOSTNAME =~ cheyenne* ) then
   setenv PATH ~/local/ch/bin:$PATH
else
   setenv PATH ~/local/dav/bin:$PATH
endif

# Settings for interactive shells
if ( $?prompt ) then
   set prompt = "%n@%m:%~"
endif
```