

Using NCAR's HPC Resources

With Focus on the New Supercomputer *Derecho*

Ben Kirk
CISL Consulting Services Group
benkirk@ucar.edu

https://arc.ucar.edu/knowledge_base_documentation



September 25, 2023



- 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



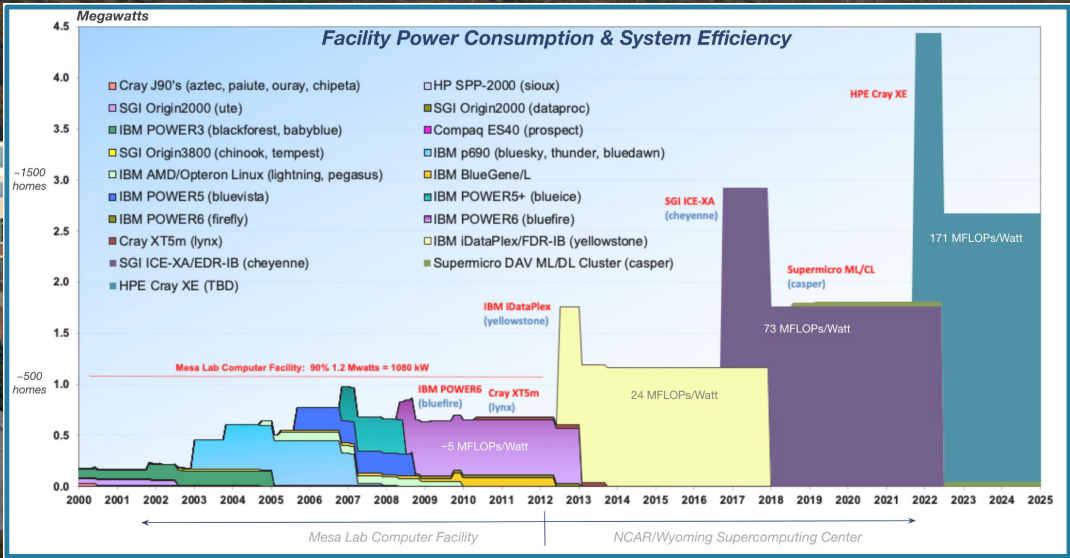
8120 Veta Drive





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 & -3 Cheyenne & Derecho supercomputers
-
- 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



For NWSC Virtual Tour visit:

<https://bit.ly/NCAR360>



Derecho Construction Virtual Tour is available for viewing at:


<https://www.thinglink.com/mediacard/1510396772545986561>



NCAR's High-Performance Computing, Data, & Analysis Resources


HPC Systems

2017 SGI/HPE
4032 Nodes, **145,152 Cores**, 313 TB total memory, **4.79 PFlop/s**
#21 Supercomputer in the world at debut



2023

mid-2023 Cray/HPE
2570 Nodes, **323,712 CPU Cores**, 680 TB total memory
3.5X performance vs Cheyenne
328 NVidia A100 GPUs provides 20% overall performance, **19.87 PFlop/s (projected)**



Data Analysis & Visualization



Casper: heterogeneous system for data analysis & viz.

- 75 High-Throughput Computing nodes
- 9 visualization nodes with accelerated graphics
- 18 dense GPU nodes for AI/ML, Code Development
- 4 nodes for Research Data processing
- 2 1.5TB large memory nodes

<http://projectpythia.org> **CISL develops specialized visualization software & services for Earth Science applications**






<https://geocat.ucar.edu>




High Performance Storage



GLADE & Campaign Storage

- 150PB long-term, online storage
- 17,464 hard drives
- 56 servers



Derecho 'scratch' Storage


- 60PB short-term storage
- Principally supports HPC jobs
- 5,088 hard drives
- 24 servers

Stratus Object Storage

- 5PB object storage system
- 588 hard drives
- 6 servers

Quasar Tape Library

- 35PB long term archival storage
- 22 IBM TS1160 tape drives
- 1774 20TB tape cartridges
- 216 hard drives
- 2PB disk cache
- 5 data mover servers





CRAY

HPCCray

HPCCray

HPCCray

HPCCray

HPCCray

HPCCray

HPCCray

HPCCray

Derecho

HPE/Cray EX Supercomputer

- 3rd supercomputing system deployed at NWSA
- Entered production June 2023 (early users), general availability August 2023

2488 CPU-only Compute nodes (318,464 total cores)

- Dual socket, 64 cores per socket
- AMD EPYC™ 7763 (Milan) processors
- 637 TB total system memory, 256GB/node, DDR4
- Cray Slingshot 11 Dragonfly Network, 1X 200Gb/s network interface

82 GPU Compute nodes (328 total A100 general purpose GPUs)

- Single socket, 64 cores per socket
- AMD EPYC™ 7763 (Milan) processor
- 4X NVIDIA 1.41 GHz A100 Tensor Core GPUs per node with 40GB memory and 600 GB/s NVIDIA NVLink GPU interconnect
- 42 TB total system memory, 512GB/node, DDR4
- Cray Slingshot 11 Dragonfly Network, 4X 200Gb/s network interface

8 login nodes

Documentation Home Page:

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



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

Planned system lifetime through December 2023

Documentation Home Page:

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





Casper

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

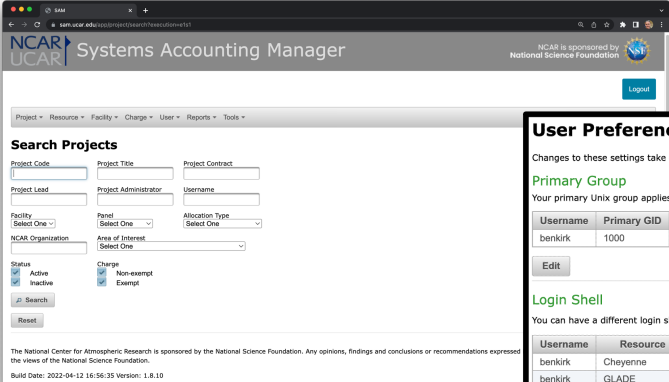
- **75 high-throughput computing (HTC)** nodes across 2 generations of hardware for small computing tasks using 1 or 2 CPUs.
 - Typical HTC nodes have 384 GB system memory.
- **2 large memory nodes** have 1.5 TB system memory.
- **9 nodes for data analysis and visualization jobs.**
 - These nodes include a single NVIDIA Quadro GP100 16 GB GPU and 384 GB system memory.
- **18 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.
 - 8 of these nodes feature 8 NVIDIA Ampere A100 80 GB GPUs and 1 TB system memory.
- **4 nodes** are reserved for **Research Data Archive** workflows.

Users must be very explicit when requesting required resources because Casper is highly heterogeneous (more later).

Documentation Home Page:

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

SAM – Systems Accounting Manager



User Preferences

Changes to these settings take effect next business day.

Primary Group

Your primary Unix group applies to all CISL resources.

Username	Primary GID	Primary Group Name
benkirk	1000	ncar

Edit

Login Shell

You can have a different login shell on each resource accessible to you.

Username	Resource	Shell
benkirk	Cheyenne	bash
benkirk	GLADE	bash
benkirk	HPC_Futures_Lab	bash
benkirk	Laramie	bash

Edit

- 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 Derecho, Cheyenne, or Casper

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

```
ssh -Y username@derecho.hpc.ucar.edu
```

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

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

(The `-Y` is optional, allows graphics forwarding when desired)

- Use password+Duo for two-factor authentication
- You will be placed on a login node
 - Derecho - 8 login nodes: `derecho[1-8]`
 - Cheyenne - 6 login nodes: `cheyenne[1-6]`
 - Casper - 2 login nodes: `casper-login[1,2]`

***Derecho* has scheduled maintenance downtime the first Tuesday of each month**



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: Wed Sep 20 10:19:57 2023 from 71.229.188.226
*****
*                               Welcome to Casper - September 20, 2023
*                               *****
*                               Today in the Daily Bulletin (arc.ucar.edu)
*
*                               - Reminder: Migrate Your Work from Cheyenne to Derecho
*                               - NCAR Nominated for three HPCwire Readers' Choice Awards
*
Documentation:      https://ncar.pub/CISL-user-documentation
Key module commands:  module list, module avail, module spider, module help
Consulting Services: https://rchelp.ucar.edu
CISL Help:         cislhelp@ucar.edu -- 303-497-2000
-----
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	Backups	Technology	Uses
Home /glade/u/home/\${USER}	50 GB	Yes	IBM Spectrum Scale	Users' settings, source code, scripts
Work /glade/work/\${USER}	1 TB*	No	IBM Spectrum Scale	Compiled codes, models
Scratch (Derecho) /glade/derecho/scratch/\${USER}	30 TB	NO!! Purged!	Cray ClusterStor Lustre	Run Directories, Temporary outputs Purged at 180 days
Scratch (Cheyenne) /glade/cheyenne/scratch/\${USER}	10 TB	NO!! Purged!	IBM Spectrum Scale	Run Directories, Temporary outputs Purged at 120 days
Campaign Storage /glade/campaign/<PROJECTS>	<i>Allocated, Project Specific</i>	No	IBM Spectrum Scale	Project Spaces, Long-term Curated Data Sets. Automated file compression.

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

GLADE: \$HOME: snapls

```
# list $HOME directory snapshots
```

```
$ cd $HOME && snapls | grep snapshots | sort -r
```

```
... Sep 20 10:43 /glade/u/home/.snapshots/20230920-130001/benkirk
... Sep 20 10:43 /glade/u/home/.snapshots/20230920-120001/benkirk
... Sep 20 10:43 /glade/u/home/.snapshots/20230920-110001/benkirk
... Sep 20 08:35 /glade/u/home/.snapshots/20230920-090002/benkirk
... Sep 19 15:09 /glade/u/home/.snapshots/20230920-060001/benkirk
... Sep 19 15:09 /glade/u/home/.snapshots/20230920-030001/benkirk
... Sep 19 15:09 /glade/u/home/.snapshots/20230920-000001/benkirk
... Sep 19 15:09 /glade/u/home/.snapshots/20230919-180002/benkirk
... Sep 19 11:08 /glade/u/home/.snapshots/20230919-120001/benkirk
... Sep 18 08:34 /glade/u/home/.snapshots/20230919-000001/benkirk
... Sep 14 15:12 /glade/u/home/.snapshots/20230918-000001/benkirk
... Sep 14 15:12 /glade/u/home/.snapshots/20230917-000001/benkirk
... Sep 14 15:12 /glade/u/home/.snapshots/20230916-000001/benkirk
... Sep 14 15:12 /glade/u/home/.snapshots/20230915-000001/benkirk
... Sep 13 19:38 /glade/u/home/.snapshots/20230914-000001/benkirk
```

```
# peruse a snapshot ...
```

```
$ ls /glade/u/home/.snapshots/20230920-000001/benkirk
```

GLADE: \$HOME: snapls

```
# list $HOME directory snapshots
```

```
$ cd $HOME && snapls | grep snapshots | sort -r
```

```
... Sep 20 10:43 /glade/u/home/.snapshots/20230920-130001/benkirk
... Sep 20 10:43 /glade/u/home/.snapshots/20230920-120001/benkirk
... Sep 20 10:43 /glade/u/home/.snapshots/20230920-110001/benkirk
... Sep 20 08:35 /glade/u/home/.snapshots/20230920-090002/benkirk
... Sep 19 15:09 /glade/u/home/.snapshots/20230920-060001/benkirk
... Sep 19 15:09 /glade/u/home/.snapshots/20230920-030001/benkirk
... Sep 19 15:09 /glade/u/home/.snapshots/20230920-000001/benkirk
... Sep 19 15:09 /glade/u/home/.snapshots/20230919-180002/benkirk
... Sep 19 11:08 /glade/u/home/.snapshots/20230919-120001/benkirk
... Sep 18 08:34 /glade/u/home/.snapshots/20230919-000001/benkirk
... Sep 14 15:12 /glade/u/home/.snapshots/20230918-000001/benkirk
... Sep 14 15:12 /glade/u/home/.snapshots/20230917-000001/benkirk
... Sep 14 15:12 /glade/u/home/.snapshots/20230916-000001/benkirk
... Sep 14 15:12 /glade/u/home/.snapshots/20230915-000001/benkirk
... Sep 13 19:38 /glade/u/home/.snapshots/20230914-000001/benkirk
```

```
# peruse a snapshot ...
```

```
$ ls /glade/u/home/.snapshots/20230920-000001/benkirk
```

GLADE: gladequota

```
$ gladequota
```

```
Current GLADE space usage: benkirk
```

Space	Used	Quota	% Full	# Files
/glade/derecho/scratch/benkirk	1.13 PiB	1.50 PiB	75.47 %	4671187
/glade/scratch/benkirk	554.39 GiB	10.00 TiB	5.41 %	895051
/glade/work/benkirk	719.42 GiB	1.00 TiB	70.26 %	5946730
/glade/u/home/benkirk	67.69 GiB	100.00 GiB	67.69 %	20705
/glade/u/sampleddata	51.82 GiB	1.00 TiB	5.06 %	105
/glade/p/cesm	1.11 PiB	1.17 PiB	96.33 %	10314948
/glade/u/apps	1.98 TiB	10.00 TiB	19.76 %	19932245
Campaign: benkirk (user total)	2.81 PiB	n/a	n/a	31972844
/glade/campaign/cgd/cesm	1.88 PiB	2.00 PiB	93.99 %	5898982
/glade/campaign/cesm	9.53 PiB	18.50 PiB	51.51 %	44306409
/glade/campaign/cisl/csg	7.94 TiB	23.00 TiB	34.54 %	2859

(Campaign usage as of: Wed Sep 20 10:29:40 MDT 2023)

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
- As of July 2023, accessible from all HPC resources (including, now, *Cheyenne*).
- Files automatically compressed >6 months after creation; transparently to 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, GCC, NVIDIA, Cray)
- Linaro Forge Debuggers / Performance Tools (DDT, MAP)
- MPI Libraries (Cray-MPICH, 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://rhelp.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 Derecho, 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
```

```
--- Compilers and Core Software ---
```

```
atp/3.14.18                gcc/12.2.0                nvhpc/23.1 (D)
atp/3.15.0 (D)             gdb4hpc/4.14.7           nvhpc/23.5
cce/15.0.1                 go/1.20.3                papi/7.0.0.1
cdo/2.1.1                  grads/2.2.1              pcre/8.45
charliecloud/0.32         idl/8.9.0                peak-memusage/3.0.1
cmake/3.26.3              intel-classic/2023.0.0   perftools-base/23.03.0
conda/latest              intel-oneapi/2023.0.0    perl/5.36.0
cray-ccdb/4.12.13        intel/2023.0.0 (L)       podman/4.3.1
cray-dyninst/12.1.1      linaro-forge/23.0        sanitizers4hpc/1.0.4
cray-mrnet/5.0.4         matlab/R2023a            texlive/20220321
```

```
...
```

```
--- Compiler-dependent Software - [oneapi] ---
```

```
cray-libsci/23.02.1.1     geos/3.9.1                mkl/2023.0.0
cray-mpich/8.1.25 (L)    hdf/4.2.15                proj/8.2.1
eccodes/2.25.0           hdf5/1.12.2 (L)          ncarcompilers/1.0.0 (L)
fftw/3.3.10              intel-mpi/2021.8.0        netcdf/4.9.2 (L)
```

```
--- MPI-dependent Software - [oneapi + cray-mpich] ---
```

```
darshan-runtime/3.4.2    gptl/8.1.1                parallelio/1.10.1
esmf/8.4.2               hdf5-mpi/1.12.2           parallelio/2.5.10
esmf/8.5.0 (D)          netcdf-mpi/4.9.2          parallelio/2.6.0
fftw-mpi/3.3.10         parallel-netcdf/1.12.3    parallelio/2.6.1
```

Modules (*Derecho* Demonstration)

```
$ module list
```

```
Currently Loaded Modules:
```

```
1) ncarenv/23.06 (S)   3) intel/2023.0.0       5) cray-mpich/8.1.25   7) netcdf/4.9.2  
2) craype/2.7.20      4) ncarcompilers/1.0.0  6) hdf5/1.12.2
```

```
$ which mpicc && echo $NETCDF && echo $NCAR_BUILD_ENV
```

```
/glade/u/apps/derecho/23.06/spack/opt/spack/ncarcompilers/1.0.0/oneapi/2023.0.0/ec7b/bin/mpi/m  
/glade/u/apps/derecho/23.06/spack/opt/spack/netcdf/4.9.2/oneapi/2023.0.0/iijr  
derecho-oneapi-19-cray-mpich-8.1.25
```

Modules (*Derecho* Demonstration)

```
$ module swap intel gcc/12.2.0
```

Due to MODULEPATH changes, the following have been reloaded:

1) cray-mpich/8.1.25 2) hdf5/1.12.2 3) ncarcompilers/1.0.0 4) netcdf/4.9.2

Modules (*Derecho* Demonstration)

```
$ module swap intel gcc/12.2.0
```

Due to MODULEPATH changes, the following have been reloaded:

```
1) cray-mpich/8.1.25    2) hdf5/1.12.2    3) ncarcompilers/1.0.0    4) netcdf/4.9.2
```

```
$ which mpicc && echo $NETCDF && echo $NCAR_BUILD_ENV
```

```
/glade/u/apps/derecho/23.06/spack/opt/spack/ncarc compilers/1.0.0/gcc/12.2.0/gj6c/bin/mpicc  
/glade/u/apps/derecho/23.06/spack/opt/spack/netcdf/4.9.2/gcc/12.2.0/ok4t  
derecho-gcc-12.2.0-cray-mpich-8.1.25
```

```
$ module list
```

Currently Loaded Modules:

```
1) ncarenv/23.06 (S)    3) gcc/12.2.0    5) cray-mpich/8.1.25    7) netcdf/4.9.2  
2) craype/2.7.20    4) ncarcompilers/1.0.0    6) hdf5/1.12.2
```









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 mk1
module save default
```
- You can make multiple named sets and load them using `module restore <set>`

Considerations when compiling software

- Use **ncarccompilers** 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**
 - Derecho, Cheyenne, and Casper have different CPUs, operating systems, and default MPI implementations





Running Jobs

Leigh Orin
l.Orin@ucar.edu

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



Example PBS batch job scripts (Cheyenne)

```
$ cat basic_mpi_cheyenne.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}/tmp && mkdir -p ${TMPDIR}

module load intel/19.1.1 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:

- Derecho: https://arc.ucar.edu/knowledge_base/131596447
- Cheyenne: https://arc.ucar.edu/knowledge_base/72581486
- Casper: https://arc.ucar.edu/knowledge_base/72581394

Example PBS batch job scripts (Derecho, CPU)

```
$ cat basic_mpi_derecho.pbs
#!/bin/bash
#PBS -N hello_pbs
#PBS -o pbsjob.log
#PBS -A <project_code>
#PBS -j oe
#PBS -k eod
#PBS -q main
#PBS -l walltime=00:05:00
### Select 2 nodes with 128 CPUs each for a total of 256 MPI processes
#PBS -l select=2:ncpus=128:mpiprocs=128:ompthreads=1

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

module load intel/2023.0.0 cray-mpich/8.1.25 && module list

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

### Run Cray-MPICH MPI Program
mpiexec -n 256 -ppn 128 ./hello_world
```

More Examples:

- Derecho: https://arc.ucar.edu/knowledge_base/131596447
- Cheyenne: https://arc.ucar.edu/knowledge_base/72581486
- Casper: https://arc.ucar.edu/knowledge_base/72581394

Example PBS batch job scripts (Derecho, GPU)

```
$ cat basic_mpi_derecho_gpu.pbs
#!/bin/bash
#PBS -A <project_code>
#PBS -N gpu_job
#PBS -q main
#PBS -j oe
#PBS -k eod
#PBS -l walltime=01:00:00
#PBS -l select=2:ncpus=64:mpiprocs=4:ngpus=4

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

### load necessary module environment
module purge
module load ncarenv nvhpc cuda cray-mpich

### Run application with efficient mapping to GPUs and NICs
export MPICH_OFI_NIC_POLICY=GPU
mpiexec -n 8 -ppn 4 set_gpu_rank ./executable_name
```

More Examples:

- Derecho: https://arc.ucar.edu/knowledge_base/131596447
- 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.chadmi1.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 | grep "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

Remote Desktop & Graphical Programs: VNC

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

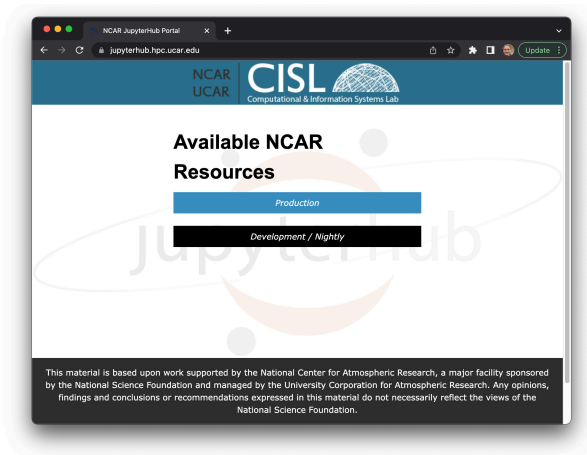
```
# launch 'vncmngnr' on casper, matched with a VNC client locally  
casper$ vncmgr
```



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

JupyterHub

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



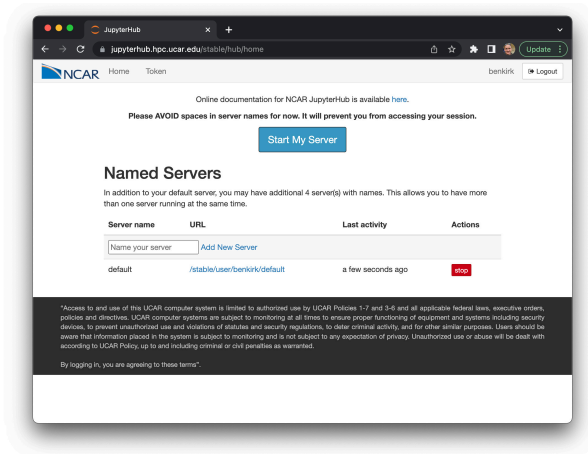
<https://jupyterhub.hpc.ucar.edu/>

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



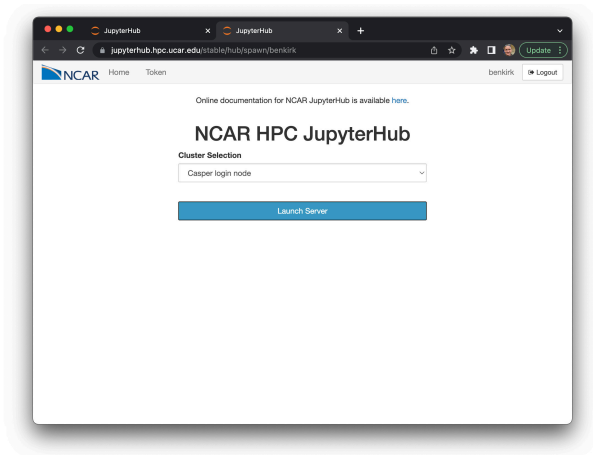
<https://jupyterhub.hpc.ucar.edu/>

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



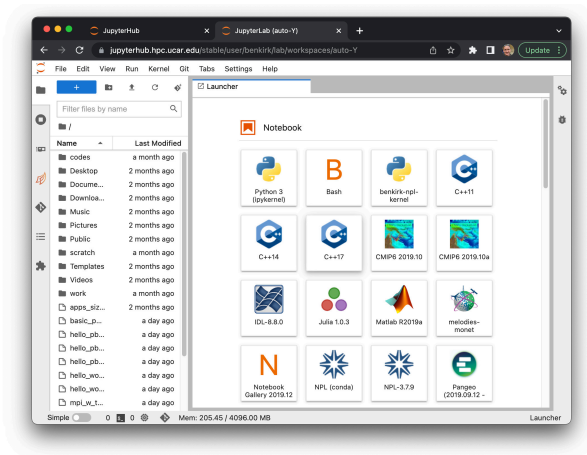
<https://jupyterhub.hpc.ucar.edu/>

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



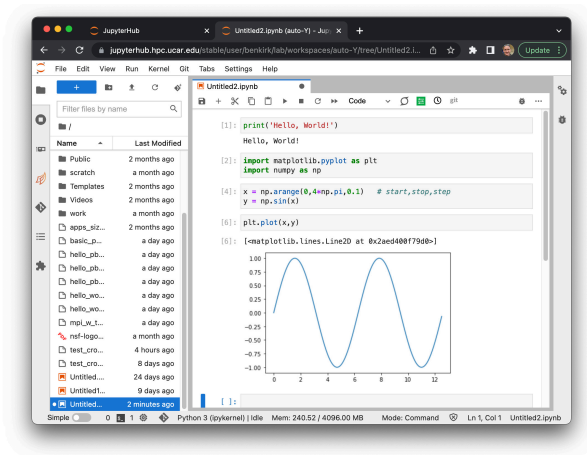
<https://jupyterhub.hpc.ucar.edu/>

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



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

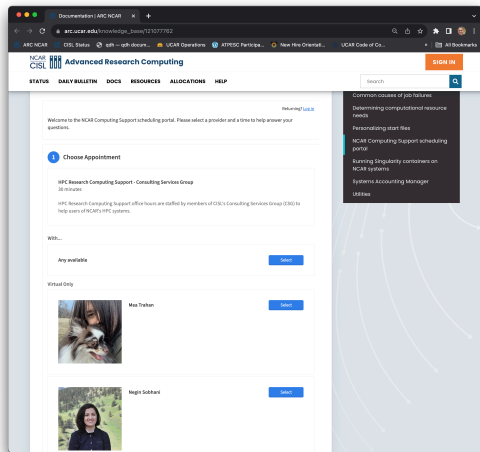
When submitting a support ticket please include as much detail as possible to enable quicker resolution:

- Resource name (Derecho, 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>

Virtual Consulting Support

- CISL's Consulting Services Group (CSG) is pleased to announce the availability of virtual consulting services by appointment: **NCAR Computing Support scheduling portal**
- Intent is to augment the primary **ticket support mechanism** with a scheduled, one-on-one support model:
 - “Next Availability” scheduling available for general queries, or
 - Schedule a particular consultant, intended as a follow-on to an existing ticket
- Integrated with Google Calendar & Meet, ideal for problems that benefit from screen sharing.





wiki

<https://bit.ly/3Bco4Wh>



Slack

<http://ncarhpcusergroup.slack.com>

Check out
#derecho-users
#cheyenne-users
#casper-users



Slido

<http://www.slido.com/nhug>

Meet NCAR's HPC Consultants



Ben Kirk

Manager,
Consulting Services
Group

Joined NCAR in January
2022

Professional Interests

- MPI & hybrid parallel applications
- High performance storage
- Optimizing user workflows using knowledge of hardware strengths and weaknesses
- C++, Python

Bio

Ben attended the University of Texas at Austin, earning his BS and PhD in Aerospace Engineering, with a MS in Computational & Applied Mathematics. He worked at NASA's Lyndon B. Johnson Space Center for 18 years prior to joining NCAR. His background includes constructing and using HPC systems (compute and storage), particularly for fluid dynamics simulations.

Personal Interests

- Rock Climbing
- Snowboarding
- Hiking
- Music
- Reading Nonfiction

Meet NCAR's HPC Consultants



Brian Vanderwende

HPC User Support,
Consulting Services
Group

Joined NCAR in
November 2015

Professional Interests

- User environment and interface design
- Teaching and training
- Workflow optimization
- Debugging
- Performance tuning

Bio

Brian originally pursued Meteorology and Atmospheric Science at Penn State (BS) and the University of Colorado (PhD), respectively. Brian's graduate work running LES simulations of wind farms using WRF, writing analysis scripts in IDL and NCL, and a general interest in the technical side of science led him to NCAR - where he can combine his passions for earth science and software engineering.

Personal Interests

- Rock climbing (especially traditional & alpine)
- Hiking and peak-bagging
- Skiing
- Sustainable living and new urbanism
- PC Games

Meet NCAR's HPC Consultants



Mea Trahan

HPC Consultant
Consulting Services
Group

Joined NCAR in 2022

Professional Interests

- Software Management and Installations
- Scientific applications of Containers
- MPI programming and parallel applications
- Python and managing python with Conda

Bio

Mea started off as a Student interested in a variety of topics at the University of Colorado Boulder. She then fell into the world of HPC when she entered a student position at CU Boulder Research Computing. She shifted over to a full time team member and left the group in Mid 2022 to join the CISL Team.

Personal Interests

- Cooking
- Painting and Digital Illustration
- MMOs and other video games

Meet NCAR's HPC Consultants



Negin Sobhani

HPC Consultant
Consulting Services
Group

Joined NCAR in 2017 as
ASP post-doc

Professional Interests

- Performance tuning and optimization
- Lowering barriers to entry for weather and climate modeling through the use of various computational technologies
- Machine Learning/AI for improving weather and climate modeling
- Best practices in computational science that promote open science and reproducible workflows
- Open source software development for geospatial data analysis
- GPU Computing

Bio

Negin began her journey at NCAR as a SIParCS student and visitor during her Ph.D., focusing on optimizing weather and air quality models. This experience ignited her passion for bridging the gap between computational and atmospheric sciences. After her ASP post-doc, she joined CGD to work on developing climate model. Combining her passions for atmospheric science and software engineering, she is focused on creating user-friendly tools and frameworks to lower barriers of entry to atmospheric sciences. Recently, she joined the CSG team as an HPC consultant, enabling her to provide even more support and assistance to users.

Personal Interests

- Backpacking in summer and hot tent camping in winter.
- Amateur astrophotography
- Woodworking
- Board games and video games
- Books on social and economical issues.

Meet NCAR's HPC Consultants



Daniel Howard

HPC Consultant I,
Consulting Services
Group

Joined NCAR Feb 2021

B.S. Applied Math, Engineering, &
Physics (AMEP), UW-Madison
M.S. Applied & Computational Math &
Stats, Notre Dame

Professional Interests

- General Purpose GPU Computing
- Computational performance of math algorithms
- Computational Fluid Dynamics methods, RBF-FD
- Open source software and scientific community development, ie hackathons
- Beyond Fortran/C++, interest in Julia, Python, and geospatial data pipelines

Member of [US-RSE.org](https://www.us-rse.org) (Research Software Engineer) and Secretary of OpenACC organization

Bio

Daniel's primary background is in applied mathematics, with intersectional interests in computational and numerical methods applied to the geosciences. Prior to NCAR, Daniel was in the Climate/Weather/Ocean division of the Naval Research Lab at Stennis Space Center working on GPU development of WaveWatchIII and other next gen architecture projects. He has been engaged with the HPC community since grad school, when he learned of its importance to modern science, including an internship in Kobe, Japan in the parallel algorithms research team at RIKEN's K Computer in 2018.

Personal Interests

- Biking, hiking, rock climbing, and calisthenics
- Plant based cooking, e.g. Indian cuisine
- Learning about urban planning
 - Highlighting opportunities for people centric places, mixed use zoning, biking infrastructure, and replacing car dependent design
 - Member of [Urban Environmentalists](https://www.urbanenvironmentalists.org)
- Local science policy and communication, as through Engineers and Scientists Acting Locally ([ESAL.us](https://www.esal.us)) and local Sierra Club chapter
- Narrative focused indie, RPG, & JRPG video games

Meet NCAR's HPC Consultants



Rory Kelly

HPC Consultant, Consulting Services Group

Joined the Computational Science Section of the Scientific Computing Division at NCAR in 2001

Bio

Rory attended CU Boulder where he studied mathematics and physics. He worked in a high energy physics research group performing monte carlo simulations of mSUGRA class supersymmetric particle decays to help design detectors for proposed linear collider. The collider was never built and the entire parameter space he studied has subsequently been excluded, but accidentally learning Fortran led to working at NCAR. Rory abandoned grad school to race bicycles, and then abandoned bicycles to race skis, and then abandoned skis to hang out with a baby, who he will inevitably race when she's old enough.

Professional Interests

- Performance measurement and performance optimization, benchmarking.
- Finding the root cause of initially baffling problems on HPC machines.
- Programming languages generally, Haskell, Julia, and C, specifically
- Making code work on weird hardware
- Writing small benchmarks to explore specific aspects of hardware.

Personal Interests

- Running in mountains, climbing in mountains, skiing in mountains, biking in mountains, looking at mountains.
- Maps and route planning.
- Hanging out with animals and babies.
- Sharpening knives, axes, scissors.
- Sailing, both leisure and performance.
- Racing things

Meet NCAR's HPC Consultants



Brett Neuman

HPC Consultant, Consulting Services Group

Joined NCAR in May 2022

Bio

Brett started as a student assistant for CSG in 2022 and transitioned to full-time in June of 2023. He is attending the University of Minnesota for his MS in Electrical & Computer Engineering and earned his BS in Computer Engineering from the University of Wisconsin. Before working at NCAR, Brett worked on FPGA prototyping, mixed precision designs, reproducibility, and co-design at Los Alamos National Laboratory.

Professional Interests

- GPUs, FPGAs, and weird architectures
- Profiling and benchmarking
- HW/SW Co-design
- Workflow and system tools

Personal Interests

- Mountain sports (running, climbing, hiking, skiing, ice climbing, snowy couloirs)
- Electronic designs for musical purposes
- Ice Hockey
- Reading
- Reducing single use plastic

Meet NCAR's HPC Consultants



Dick Valent

- HPC User Support in the Consulting Services Group
- Joined NCAR in 1969

Professional Interests

- Assisting NCAR cluster users
- Familiarizing users with NCAR user documentation
- Providing NCAR CSG historical context

Bio (specific to NCAR)

- 1969-1973, student appointment
- 1973-2004 full -time employment in the Consulting Services Group
- 2004-present casual, part-time employment in the Group

Personal Interests

- Vegetable gardening
- Healthy running and exercise
- Paleolithics
- Research in mathematical lattice theory

The image shows a vast mountain landscape. In the foreground, a grassy hillside slopes down towards a large, multi-story building with a modern, blocky design. The building is situated on a relatively flat area surrounded by some trees. Behind the building, the terrain rises into a dense forest of evergreen trees. In the background, towering, rugged mountains with prominent rock formations and vertical fissures rise against a clear sky. The overall scene is bright and clear, suggesting a sunny day.

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
```