

Introduction to NCAR HPC Systems

For CISL's 2021 SIParCS Cohort
and all new NCAR HPC users

Mick Coady

CISL Consulting Services Group
mickc@ucar.edu

May 24, 2021



Welcome!

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).

Welcome!

UCAR Code of Conduct

By registering to participate in this event, you agreed to:

- Adopt the values expressed in the Code of Conduct.
- Engage in respectful communication only.

That applies to verbal discussion as well as chat remarks.

Thank you!

Introduction to NCAR HPC Systems

Topics We'll Cover Today

- HPC Compute Systems
 - Cheyenne
 - Casper
 - Computing environments
- Data Storage Spaces
 - GLADE
 - Campaign Storage
 - Stratus and Quasar
- Helpful Resources

Computing Systems

- Cheyenne



Computing Systems

Cheyenne



NCAR's flagship supercomputer

- Built by SGI in 2017
- Began at #20 on the World's Top 500
 - still ranked at #60
- 6 login nodes
- 4,032 compute nodes
 - Broadwell processors
 - 145,152 Intel Xeon cores
- 313 TB total memory
 - 3,168 nodes with 64 GB/node
 - 864 nodes with 128 GB/node (large memory)

Computing Systems

Cheyenne



- PBS Pro Resource manager
 - The batch job scheduler
- OS: SUSE Enterprise Linux
- More than 3,000 users
 - 1700 users active in last 12 months
 - Average of 765 users each month
 - More than 707,000 jobs in April!
- Cheyenne Home Page
<https://www2.cisl.ucar.edu/resources/computational-systems/cheyenne>

Computing Systems

Casper



Computing Systems

Casper Overview



- Traditionally the **Data Analysis and Visualization (DAV) System**
- Recently expanded for **High Throughput Computing (HTC)**
- Built by CISL
- OS: CentOS
- Uses the PBS Pro Resource manager
(Same batch job scheduler as Cheyenne)
- Casper Home Page:
<https://www2.cisl.ucar.edu/resources/computational-systems/casper>

Computing Systems

Casper Overview (cont)



- Total 100 compute nodes
 - Heterogeneous: mix of node types
 - Skylake and Cascade Lake
 - CPU-only and GPU nodes
- 2 login nodes
- 64 **H**igh **T**hroughput **C**omputing (HTC) Nodes *<new!>*
 - 62 nodes - 384 GB memory
 - 2 nodes - 1.5 TB memory
- 4 nodes reserved for Research Data Archive (RDA)

Computing Systems

Casper Overview (cont)



- 19 GPU Nodes
 - 9 nodes with 1 NVIDIA GP100, 384 GB memory
 - Good for lightweight visualization tasks
 - 10 nodes with NVIDIA Tesla V100s for GPGPU and ML/DL
 - 4 nodes with 4 NVIDIA Tesla V100s, 768 GB memory
 - 6 nodes with 8 NVIDIA Tesla V100s, 1152 GB memory

Computing Systems

Logging into Cheyenne and Casper

- Use ssh along with your user name to log in
 - `ssh -Y -l <your_username> cheyenne.ucar.edu`
 - `ssh -Y -l <your_username> casper.ucar.edu`
- Use Duo for authentication
- You will be placed on one of the systems' login nodes
 - Round robin assignments
- Cheyenne and Casper use separate project allocations.
 - Check on your project(s) using SAM (more about that later)

Computing Systems

Cheyenne and Casper Login Nodes Usage

Be mindful of your usage on the login nodes!

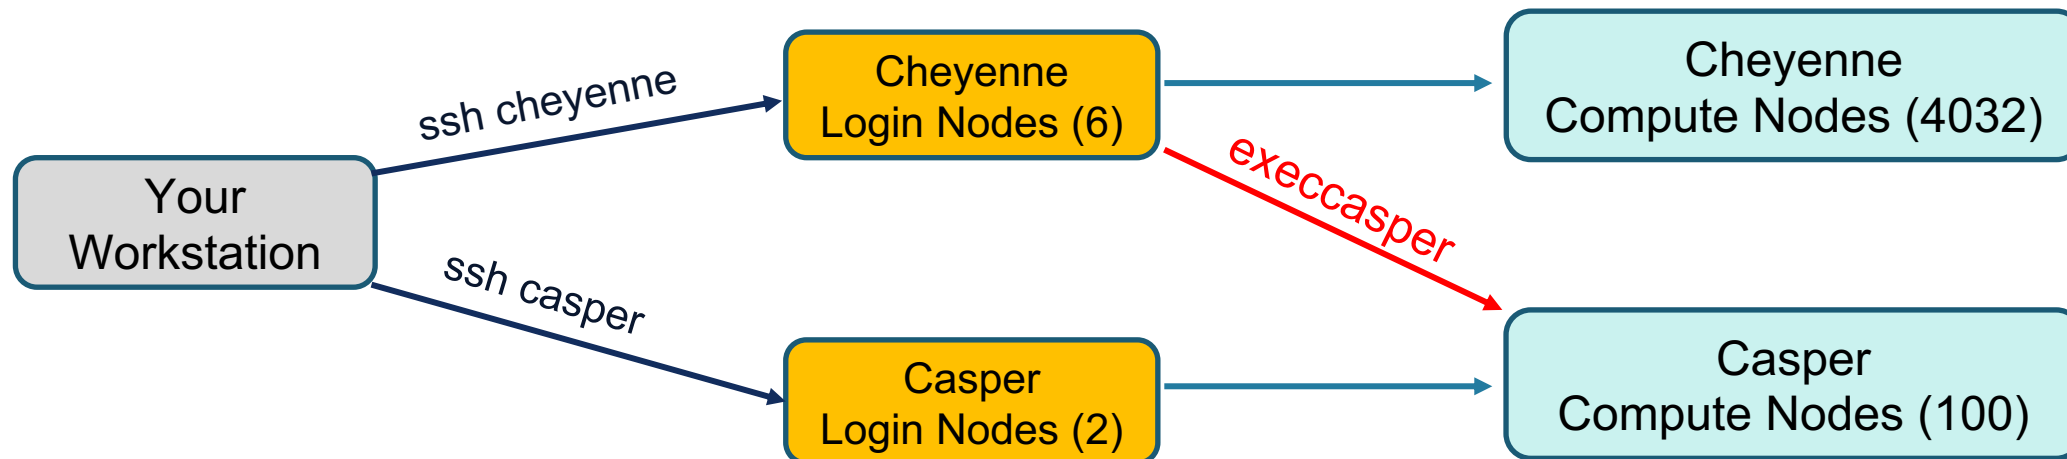
- Your activities co-exist with those of other users
- CPUs and memory are shared on the login nodes
- Limit your usage to:
 - Reading, writing, and editing text/code
 - Compiling small-ish programs (*Compile your code on the system where you intend to run*)
 - Interacting with PBS
 - Data transfers
- Programs that use excessive resources on the login nodes will be terminated 😞

Computing Systems

Cheyenne and Casper Login Nodes Usage

Be mindful of your usage on the login nodes!

- Most HPC tasks require too many resources to run on login nodes
- Run those tasks on Cheyenne or Casper compute nodes using PBS



Computing Systems

Cheyenne and Casper Software

CISL provides and maintains many system software packages

- Compilers (Intel, GNU, NVIDIA)
 - Debuggers / Performance Tools (DDT, MAP)
 - MPI Libraries (MPT, Intel MPI, OpenMPI)
 - IO Libraries (NetCDF, PNetCDF, HDF5)
 - Analysis Languages (Python, Julia, R, IDL, Matlab)
 - More ...
-
- CISL installed software is provided as [modules](#)
 - Modules prevent loading incompatible software into your environment

Computing Systems

Cheyenne and Casper Software

Environment Modules

- Cheyenne and Casper have different CPUs and operating systems so ...
- Cheyenne and Casper have independent sets of modules (but they might look similar!)
- We **strongly** recommend you build code on the system where it will run.
 - Actually, it's a necessity!

Computing Systems

Cheyenne and Casper Software

Default Environment Modules

Cheyenne

- ncarenv
- Intel
- ncarcompilers
- **mpt**
- netcdf

Casper

- ncarenv
- Intel
- ncarcompilers
- **openmpi**
- netcdf

Intel is the default compiler on both systems

Environment Modules home page:

<https://www2.cisl.ucar.edu/resources/computational-systems/cheyenne/user-environment/environment-modules>

Computing Systems

Cheyenne and Casper Software

Using Modules – some commonly used commands

- **module list** – List the modules that are loaded
- **module load <software>** – Load the default version of a software package or load a specific version
- **module unload < software>** – Unload the specified software package
- **module purge** – Unload all modules
- **module save <name>** – create/save the current set of modules
- **module restore <name>** – load a saved set of modules
- **module avail** – list all currently loadable modules
- **module spider** – list all modules

Computing Systems

Cheyenne and Casper Software

Using Modules – shortcut aliases

- **ml = module list**
- **ml <software> = module load <software>**
- **ml unload <software> = module unload < software>**
- **ml purge = module purge**
- **ml save <name> = module save <name>**
- **ml av = module avail**
- **ml spider = module spider**

Computing Systems

Cheyenne and Casper Software

Using Modules – examples

```
mickc@cheyenne3:~--> module list
```

```
Currently Loaded Modules:
```

```
1) ncarenv/1.3    2) intel/19.0.5    3) ncarcompilers/0.5.0    4) mpt/2.22    5) netcdf/4.7.4
```

Computing Systems

Cheyenne and Casper Software

Using Modules – examples

```
mickc@cheyenne3:~--> module list
```

```
Currently Loaded Modules:
```

```
1) ncarenv/1.3    2) intel/19.0.5    3) ncarcompilers/0.5.0    4) mpt/2.22    5) netcdf/4.7.4
```

```
mickc@cheyenne3:~--> module load gnu
```

```
Lmod is automatically replacing "intel/19.0.5" with "gnu/9.1.0".
```

```
Due to MODULEPATH changes, the following have been reloaded:
```

```
1) mpt/2.22    2) ncarcompilers/0.5.0    3) netcdf/4.7.4
```

Computing Systems

Cheyenne and Casper Software

Using Modules – examples

```
mickc@cheyenne3:~--> module list
```

```
Currently Loaded Modules:
```

```
1) ncarenv/1.3 2) intel/19.0.5 3) ncarcompilers/0.5.0 4) mpt/2.22 5) netcdf/4.7.4
```

```
mickc@cheyenne3:~--> module load gnu
```

```
Lmod is automatically replacing "intel/19.0.5" with "gnu/9.1.0".
```

```
Due to MODULEPATH changes, the following have been reloaded:
```

```
1) mpt/2.22 2) ncarcompilers/0.5.0 3) netcdf/4.7.4
```

```
mickc@cheyenne3:~--> module list
```

```
Currently Loaded Modules:
```

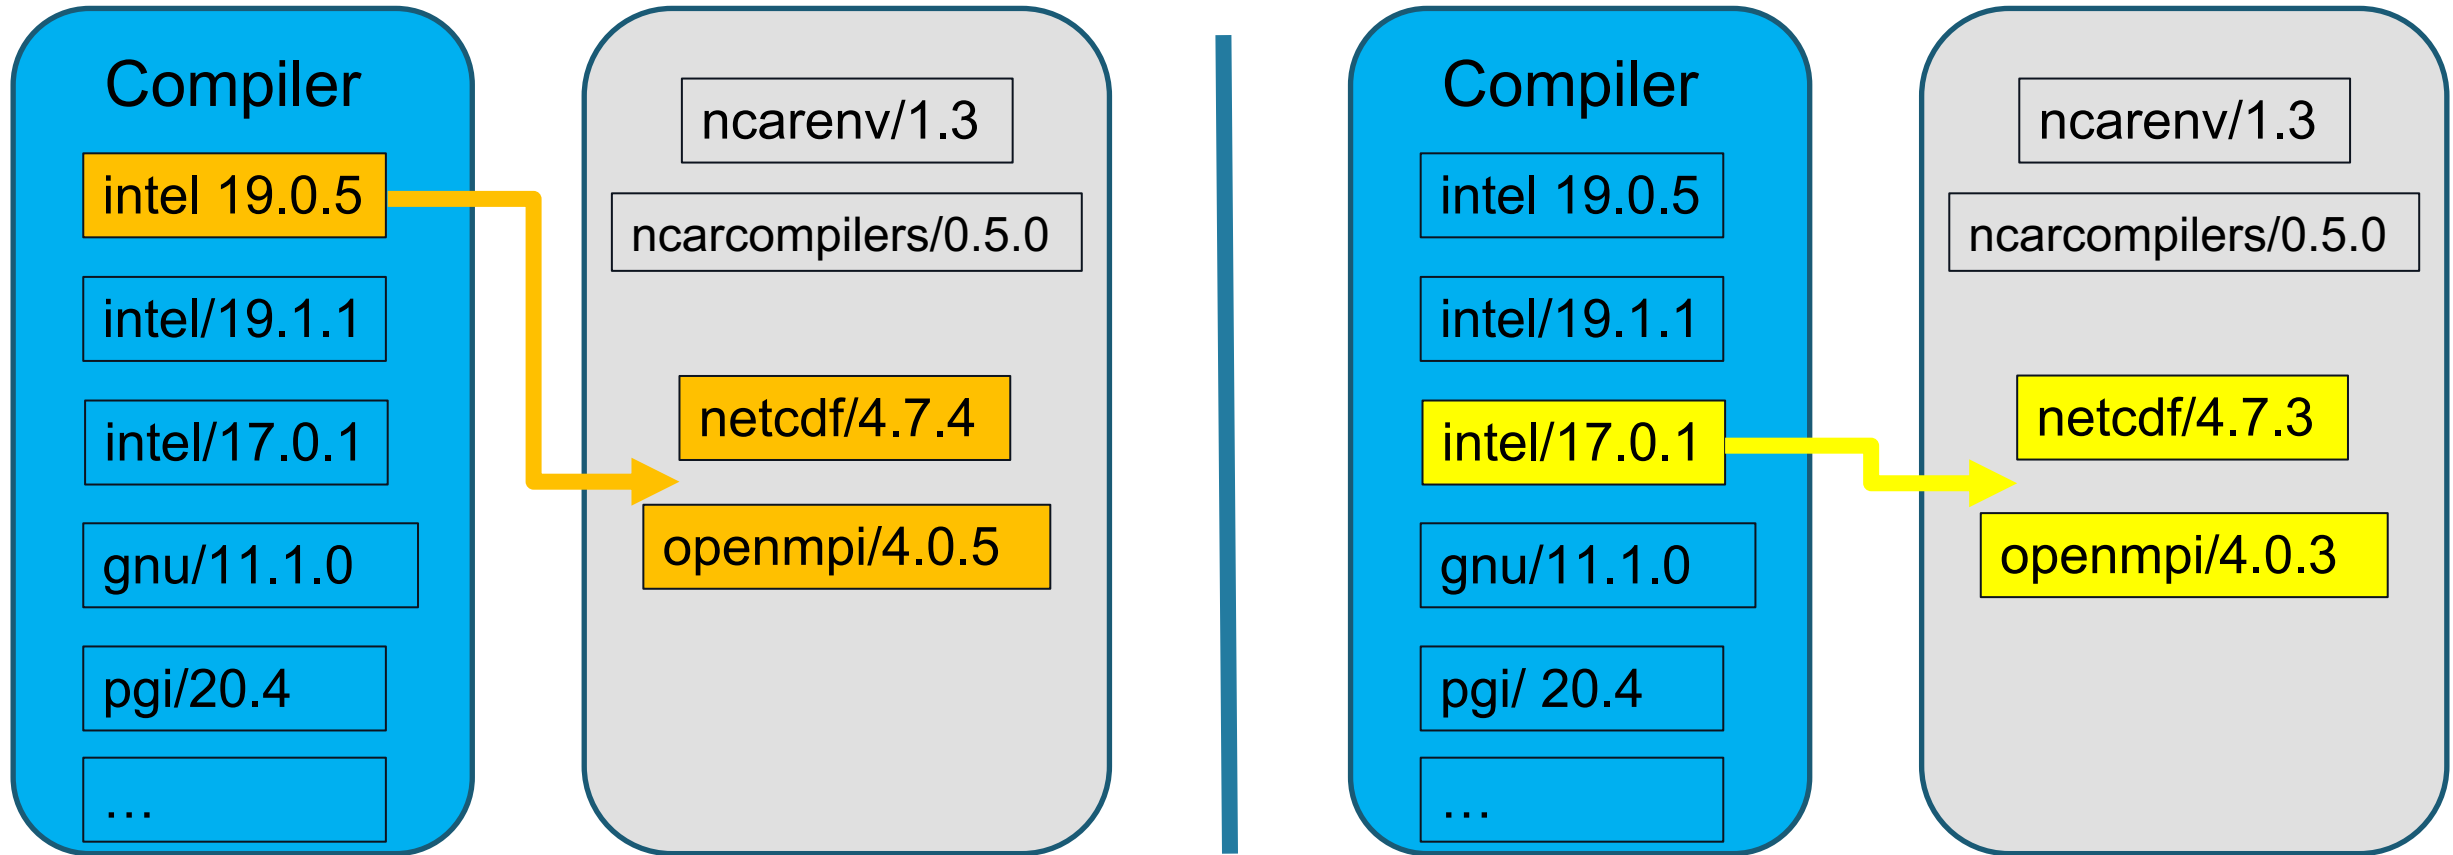
```
1) ncarenv/1.3 2) gnu/9.1.0 3) ncarcompilers/0.5.0 4) netcdf/4.7.4 5) mpt/2.22
```

Module
Magic!

Computing Systems

Cheyenne and Casper Software

Default modules \longrightarrow `module load intel/17.0.1` \longrightarrow Loaded modules



Computing Systems

Cheyenne and Casper Software

Using the PBS Resource Manager

- `qsub <PBS script>` submit a batch job
- `qstat <job_ID>` query a job's status
- `qstat -u $USER` query all jobs belonging to \$USER
- `qdel <job_ID>` kill a job

PBS commands man pages are very good sources of information.
e.g. `man qstat`

CISL scripts available on Cheyenne:

- `qinteractive` - run an interactive job on a shared node (one hour max)
- `qcmd` - run a resource-intensive task on a compute node
- `execcasper` - launch an interactive session on a Casper compute node

CISL documentation:

<https://www2.cisl.ucar.edu/resources/computational-systems/cheyenne/running-jobs/submitting-jobs-pbs>

PBS documentation:

<https://www.altair.com/pdfs/pbsworks/PBSUserGuide2021.1.pdf>

Computing Systems

Cheyenne and Casper Software

Submitting Jobs – example Cheyenne script, [example.pbs](#)

```
#!/bin/tcsh
#PBS -N job_name
#PBS -A project_code
#PBS -l walltime=01:00:00
#PBS -q queue_name
#PBS -j oe
#PBS -k eod
#PBS -m abe
#PBS -M your_email_address
#PBS -l select=1:ncpus=36:mpiprocs=36

### Set TMPDIR as recommended
setenv TMPDIR /glade/scratch/$USER/temp
mkdir -p $TMPDIR

### Run the executable
mpiexec_mpt ./executable_name.exe
```

Substitute with your own job name, project code, queue specification, email address, and so on where indicated.

Request one node and all 36 CPUs

Computing Systems

Cheyenne and Casper Software

Submitting Jobs – example Cheyenne script, [example.pbs](#)

```
#!/bin/tcsh
#PBS -N example_jpb
#PBS -A SCSG0001
#PBS -l walltime=01:00:00
#PBS -q regular
#PBS -j oe
#PBS -k eod
#PBS -m abe
#PBS -M mickc@ucar.edu
#PBS -l select=1:ncpus=36:mpiprocs=36

### Set TMPDIR as recommended
setenv TMPDIR /glade/scratch/$USER/temp
mkdir -p $TMPDIR

### Run the executable
mpiexec_mpt ./hello_world.exe
```

```
mickc@cheyenne3:~--> qsub example.pbs
```

Introduction to NCAR HPC Systems

Data Storage Systems

- GLADE
- Campaign Storage
- Quasar
- Stratus

Data Storage Systems

GLADE

- **GLADE – GLobaly Accessible Data Environment**
- High performance, disk storage space
- Where most user's active files are stored
- **GLADE home page:**
<https://www2.cisl.ucar.edu/resources/storage-and-file-systems/glade-file-spaces>

Data Storage Systems

GLADE

File Space	Path	Quota	Backed Up?	Uses
Home	/glade/u/home/\$USER	25 GB	Yes	code, scripts, small files
Work	/glade/work/\$USER	1 TB	<u>No</u>	compiled code, models
Scratch	/glade/scratch/\$USER	10 TB	<u>No</u> /Purged*	temp output, files
Project	/glade/p/<entity>/project	n/a	<u>No</u>	Project space allocations

* Current scratch space purge policy: 120 days
Subject to change.

Data Storage Systems

GLADE

Monitor your usage with the “[gladequota](#)” command

```
mickc@cheyenne3:~--> gladequota
Current GLADE space usage: mickc
```

Space	Used	Quota	% Full	# Files
/glade/scratch/mickc	118.45 GB	10.00 TB	1.16 %	7577
/glade/work/mickc	278.68 GB	1024.00 GB	27.21 %	73704
/glade/u/home/mickc	15.19 GB	50.00 GB	30.38 %	70359
/glade/u/cesm-scripts	250.99 GB	1024.00 GB	24.51 %	250921
/glade/p/cesm	1153.85 TB	1200.00 TB	96.15 %	12293708
/glade/u/apps	893.35 GB	1.50 TB	58.16 %	10415956
Campaign: mickc (user total)	97.89 GB	n/a	n/a	35
/glade/campaign/collections/cmip/CMIP6	3006.70 TB	4096.00 TB	73.41 %	3276869
/glade/campaign/cisl/csg	4.24 TB	23.00 TB	18.41 %	1680

Data Storage Systems

GLADE

Home directory snapshots and backups

- Snapshots are created several times each day
 - Users can recover snapshots themselves
 - Located in `/glade/u/home/.snapshots/<timestamp>/$USER`
 - Run the “`snaps`” command for the list of your available snapshots
- Backups are created several times each week
 - Contact CISL to recover a backup

Documentation:

<https://www2.cisl.ucar.edu/resources/storage-and-file-systems/glade/recovering-files-snapshots>

Data Storage Systems

Campaign Storage

- Resource for storing data on publication timescales
 - ~ 5 years
- Multiple access methods:
 - Globus (endpoint name: [NCAR Campaign Storage](#))
 - Casper and data access nodes: mounted as [/glade/campaign](#)
 - Not available from Cheyenne
- Allocated to and managed by NCAR labs and projects
- Allocated to University projects by request

Data Storage Systems

Quasar

- Cold archive tape storage
 - for storing curated data collections
- Not designed for serving active data
 - Write-once, access infrequently
- Contact CISL to request access
- Quasar home page

<https://www2.cisl.ucar.edu/resources/storage-and-file-systems/quasar-archive-data-collections>

Data Storage Systems

Stratus

- Object Storage system
- In other words, not POSIX!
- Write-Once, Read-Many (WORM)
- Long-term data storage
- Contact CISL to request an account
- Stratus home page

<https://www2.cisl.ucar.edu/resources/storage-and-file-systems/stratus-object-storage-system>

Data Storage Systems

- In-depth tutorial

https://www2.cisl.ucar.edu/user-support/training/library/ncar_storage_spaces

Helpful Resources

- CISL Help Desk
 - Call 303.497.2400
 - Create a ticket:
 - [Jira Service Desk ticketing system](#)
- CISL's world-class user documentation
- Course library
 - <https://www2.cisl.ucar.edu/user-support/training/library>
- CISL Daily Bulletin
 - For current HPC news and events
 - All the cool kids read it regularly!
- Resources Overview:
<https://www2.cisl.ucar.edu/resources/resources-overview>

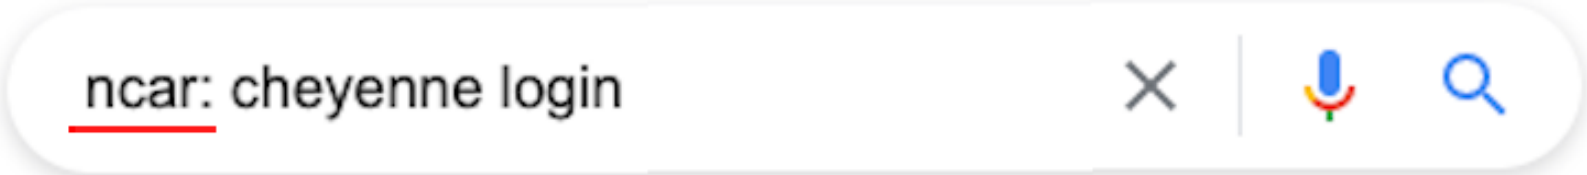
Don't be shy.
All questions and help requests
are encouraged and welcomed.

Let us know if you can't find
what you're looking for.

Helpful Resources Google Searches

Preface Google searches with keywords “ncar:” or
“ncar: cisl”

Examples:



ncar: cheyenne login



Helpful Resources

Google Searches

Preface Google searches with keywords “ncar:” or “ncar: cisl”

Examples:

Google

ncar: cheyenne login



Google

ncar: cisl cheyenne login



Helpful Resources

SAM

SAM – **S**ystems **A**ccounting **M**anager

- <https://sam.ucar.edu/>
- Requires login and Duo authentication
- Change some user settings
 - default project, shell, etc.
- Get general project information
- Check project allocation balances
- See history of jobs and charges
- Updated daily

Helpful Resources

SAM

Project ▾ Resource ▾ Facility ▾ Charge ▾ User ▾ Reports ▾ Tools ▾

Search Projects

Project Code

Project Title

Project Contract

Project Lead

Project Administrator

Username

Facility

Panel

Allocation Type

NCAR Organization

Area of Interest

Status

- Active
- Inactive

Charge

- Non-exempt
- Exempt

 Search

Thanks! Questions?

