Quick start on Cheyenne

Once you have an account and the necessary software, you can log in and run jobs on the Cheyenne supercomputer.

Logging in also gives users access to these resources, depending on their allocations:

- GLADE centralized file service
- Campaign Storage file system

To run data analysis and visualization jobs on the Casper system's nodes, follow the procedures described here. There is no need to transfer output files from Cheyenne for this since Cheyenne and Casper mount the same GLADE file systems.

> **Don’t run sudo on NCAR systems**
>
> If you need help with tasks that you think require sudo privileges, or if you aren’t sure, please contact HPC User Support before trying to run sudo yourself. The command fails when unauthorized users run it and sends a security alert to system administrators.

Page contents

Logging in on an NCAR system

To log in, start your terminal or Secure Shell client and run an **ssh** command as shown here:

```
ssh -X username@cheyenne.ucar.edu
OR
ssh -X username@casper.ucar.edu
```

Some users (particularly on Macs) need to use **-Y** instead of **-X** when calling SSH to enable X11 forwarding.

You can use this shorter command if your username for the system is the same as your username on your local computer:

```
ssh -X cheyenne.ucar.edu
OR
ssh -X casper.ucar.edu
```

After running the ssh command, you will be asked to authenticate to finish logging in.

Environment

The Cheyenne HPC system uses a Linux operating system and supports widely used shells on its login and compute nodes. Users also have several compiler choices.

Operating system

SUSE Linux.

Scheduler

Altair PBS Pro. See this detailed documentation about using PBS Pro and running jobs.

Shells

The default login shell for new Cheyenne users is **bash**. You can change the default after logging in to the Systems Accounting Manager (SAM). It may take several hours for a change you make to take effect. You can confirm which shell is set as your default by entering **echo $SHELL** on your Cheyenne command line.

Environment modules

The Cheyenne **module** utility enables users to easily load and unload compilers and compatible software packages as needed, and to create multiple customized environments for various tasks.

Here are some of the most commonly used commands. (See the Environment modules page for more details.)

module av - Show which modules are available for use with the current compiler.

module help - List switches, subcommands, and arguments for the module utility. Specify a modulefile by name for help with an individual module.

```
module help netcdf
```

module list - List the modules that are loaded.

module load - Load the default version of a software package, or load a specified version.

```
module load modulefile_name
module load modulefile_name/n.n.n
```

module spider - List all modules that exist on the system.

module swap - Unload one module and load a different one. Example:

```
module swap netcdf pnetcdf
```

module unload - Unload the specified software package.

```
module unload modulefile_name
```

Compiling

Cheyenne users have access to Intel, PGI, and GNU compilers. The **Intel compiler module** is loaded by default.

After loading the compiler module that you want to use, identify and run the appropriate compilation wrapper command from the table below. (If your script already includes one of the following generic MPI commands, there is no need to change it: mpif90, mpif77, ftn; mpicc, cc; mpiCC and CC.)

Also consider using the compiler's diagnostic flags to identify potential problems.

Compiler	Language	Commands for serial programs	Commands for programs using MPI	Flags to enable OpenMP (for serial and MPI)
Intel (default)	Fortran	ifort foo.f90	mpif90 foo.f90	-qopenmp
	C	icc foo.c	mpicc foo.c	-qopenmp
	C++	icpc foo.C	mpicxx foo.C	-qopenmp
Include these flags for best performance when you use the Intel compiler: -march=corei7 -axCORE-AVX2				
NVIDIA HPC	Fortran	nvfortran foo.f90	mpif90 foo.f90	-mp
	C	nvc foo.c	mpicc foo.c	-mp
	C++	nvc++ foo.C	mpicxx foo.C	-mp
GNU (GCC)	Fortran	gfortran foo.f90	mpif90 foo.f90	-fopenmp
	C	gcc foo.c	mpicc foo.c	-fopenmp
	C++	g++ foo.C	mpicxx foo.C	-fopenmp
PGI	Fortran	pgfortran (or pgf90 or pgf95) foo.f90	mpif90 foo.f90	-mp
	C	pgcc foo.c	mpicc foo.c	-mp
	C++	pgcpp (or pgCC) foo.C	mpicxx foo.C	-mp
The PGI compiler has become the NVIDIA HPC (nvhpc) compiler and all future versions will be released as such. PGI users should migrate to NVIDIA when possible.				

Debugging

CISL provides the ARM Forge tools, **DDT** and **MAP**, for debugging, profiling, and optimizing code in the Cheyenne environment.

Performance Reports is another ARM tool for Cheyenne users. It summarizes the performance of HPC application runs.

See Running DDT, MAP and PR jobs.

Cheyenne queues

Most of the Cheyenne batch queues are for exclusive use, and jobs are charged for all 36 cores on each node that is used. Jobs in the shared use "share" queue are charged only for the cores used.

The "regular" queue, which has a 12-hour wall-clock limit, meets most users' needs for running batch jobs.

See the table on this page for information about other queues.

Submitting jobs

Schedule your jobs to run on Cheyenne by submitting them through the PBS Pro workload management system. (To run jobs on the Casper cluster, see this documentation.)

To submit a Cheyenne *batch* job, use the **qsub** command followed by the name of your batch script file.

```
qsub script_name
```

See this page for job script examples.

To start an *interactive* job, use the **qsub** command with the necessary options but no script file.

```
qsub -I -l select=1:ncpus=36:mpiprocs=36 -l walltime=01:00:00 -q regular -A project_code
```

More detailed PBS Pro documentation

- Submitting jobs with PBS
- Scripts that you can adapt for your own jobs