

New user orientation

Are you new to NCAR and the CISE computing environment? Maybe just curious, wondering how you can use these resources in your scientific research? Either way, this is a great place to start.

We assume you already have some UNIX/Linux skills and a good idea of your scientific and computing objectives. Beyond that, here are some other essentials.

Page contents

- [What you need](#)
 - [Resources CISE provides](#)
 - [Using NCAR resources](#)
 - [How to log in](#)
-

What you need

A project and allocation. An allocation of core-hours and storage space defines the amount of resources that you can use on the systems that CISE manages. Think of an allocation as a resource budget for your research project. You can request an allocation with yourself as project lead or have a project lead add you to a project. Several types of allocation opportunities are described in our [Allocations documentation](#).

A project code. Your project is identified by an alphanumeric code. Be sure to use the appropriate project code when you submit batch jobs, store files, and do other work. This will help you and your colleagues keep track of how your resource allocation is being used.

A user account. To use CISE resources, you must have your own individual user account. Individuals who need user accounts can be designated when an allocation is requested, or the project lead can request accounts for additional users (graduate students or collaborators, for example) after a project has been awarded. See [User accounts and access](#).

Authentication. CISE provides users with a way to log in securely to supercomputing and data storage resources and others that are not available to the general public. See [Authentication and security](#) to learn about how to log in and for information on important security practices.

Your username and authentication credentials typically do not change even if you become associated with different projects and allocations over time.

Resources CISE provides

We provide world-class supercomputing, analysis, visualization, and data storage resources as well as software, data, and consulting services to support the atmospheric sciences community. All of these resources are closely interconnected.

Before you begin using them, please review the [responsibilities](#) that you accept along with the opportunity.

HPC systems

HPC systems are supercomputers that comprise many thousands of processor cores; in the case of Cheyenne, more than 145,000. This is where users develop and test their scientific parallel codes and submit batch jobs to perform simulations, perhaps with any of several NCAR [community models](#) and weather prediction programs.

Data analysis and visualization

The [Casper](#) cluster is a heterogeneous system of specialized data analysis and visualization resources as well as large-memory, multi-GPU nodes that support explorations in machine learning and deep learning. The system provides multiple types of nodes to meet users' varied needs, and it supports interactive use of scientific data-processing software such as [NCAR Command Language](#) and the [VAPOR](#) interactive 3-D visualization environment.

Data storage

The disk-based Globally Accessible Data Environment ([GLADE](#)) is accessible from any of the HPC, analysis, and visualization computer clusters that CISE manages. Each user has dedicated space on GLADE that includes a home directory, which is backed up, as well as scratch and work spaces for short-term use.

For storing project data on publication timescales, space is provided by allocation on the NCAR [Campaign Storage file system](#).

Scientific data collections

Many of our users find the data sets in our Research Data Archive and other repositories invaluable in their work. These data sets include meteorological and oceanographic observations, operational and reanalysis model outputs, and others that support atmospheric and geosciences research.

Consulting services

The Consulting Services Group provides expert advice about using our computing resources and related topics. These include programming, optimizing code, data analysis and post-processing, visualization, and data storage. See [User support](#) for how to reach a consultant.

Training

CISL provides training events, workshops, and other presentations each year. These include courses that participants can attend on-site or online, and many are recorded for reviewing at any time. Watch the [CISL Daily Bulletin](#) for announcements.

Using NCAR resources

Working in an HPC resource environment that is shared by dozens of institutions and hundreds of individual users may be quite different from your previous experience.

Here are a few additional topics to be aware of before you start. Please also see our [Best practices](#) page for some information that will help you make efficient use of your allocation.

Developing code

CISL provides many tools to help you develop and debug code for use on our supercomputing systems, which offer all of the programs, compilers, libraries, and other packages necessary for high-performance computing, analysis, and visualization.

Parallel codes are essential for successful computing in an HPC environment. If you aren't familiar with parallel programming, you may want to read [Parallel computing concepts](#) and also take advantage of some of our training opportunities to make the best use of these powerful supercomputing resources.

Submitting and running jobs

Because our supercomputing and analysis clusters are shared so widely, we employ a scheduling system to balance the workload between large and small jobs, to ensure that all members of our diverse user community have fair access, and to ensure that resources are used as productively as possible. Priorities for distributing the workload are determined by the CISL fair share policy, type of allocation, the user's choice of queues when submitting individual jobs, job size, and other factors.

Except for some small interactive processes that can run on login nodes, both interactive and batch jobs must be submitted for scheduling. Many compute jobs are simply too big to run interactively, so users submit most of these as batch jobs to run without manual intervention.

Analyzing results

We encourage all users to take advantage of our data analysis and visualization clusters to analyze the results of their HPC simulations. With the centralized GLADE file spaces, you don't have to move files from system to system to perform different tasks. When a batch job runs on the HPC system, for example, the data generated are stored on GLADE. You can then use other dedicated CISL resources to analyze and visualize the data without having to transfer files.

Transferring files

When you need to move data from GLADE or our other data storage systems to another institution for permanent storage or analysis, you can do that a number of different ways.

Users transfer files between Campaign Storage, GLADE file spaces, and remote systems with [Globus](#). Its Globus Connect Personal feature is a simple way to transfer files to and from your laptop or desktop computer.

We also provide [SCP and SFTP](#) capabilities through command line interface and Windows clients. These are best suited for transferring small numbers of small files.

Acknowledging NCAR and CISL support

Once you've conducted your work on CISL resources and are writing up the results for a journal article, presentation, or other published work, we ask that you acknowledge CISL and NCAR support for the computational aspects of the work.

Acknowledgements and citations: A requirement of all allocations and use of NCAR HPC resources managed by CISL, including the CMIP Analysis Platform and Research Data Archive, is to acknowledge NCAR and CISL support for your research and cite these resources in your publications. Our ability to identify supported scientific results helps ensure continued support from NSF and other sources for future HPC systems. See this page for how to acknowledge and cite these resources: [Acknowledging NCAR and CISL](#).

How to log in

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

```
ssh -X username@system_name.ucar.edu  
OR  
ssh -X username@system_name.hpc.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 system_name.ucar.edu  
OR  
ssh -X system_system_name.hpc.ucar.edu
```

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