

# Pleiades Front-End Usage Guidelines

## Category: Computing Policies

**Summary:** Use the Pleiades front-end systems (PFEs) and the bridge nodes for file editing, compiling, short debugging/testing sessions, and batch job submissions.

---

The PFEs and the bridge nodes are the front-end systems to Pleiades. They provide an environment that allows you to get quick turnaround while performing file editing, file transferring, compiling, short debugging/testing sessions, and batch job submission via PBS to a subset of the Pleiades compute nodes.

**WARNING:** The new Pleiades front-ends (pfe[20-27]) use the Intel Sandy Bridge processors. If you use a PGI compiler to build your executable, be aware that by default the executable is optimized for Sandy Bridge and will not necessarily execute on Harpertown, Nehalem-EP, or Westmere processors. To generate a single executable that will work on all Pleiades processor types, use the option

**-tp=penryn-64, nehalem-64, sandybridge-64** during compilation with PGI compilers. See [PGI Compilers and Tools](#) for more information.

You cannot "ssh" to the compute nodes except for the subset of nodes your PBS job is running on.

The bridge nodes are recommended for the following functions:

## Pre- and/or Post-Processing

The large amount of memory on the bridge nodes allows pre- and post-processing applications such as [Tecplot](#), [IDL](#), and [Matlab](#) to run faster than on the PFEs. Note that the bridge nodes have the same software as the PFEs. For a list of available applications, run the command `module avail`.

## File Transfers Between Pleiades and Columbia

Both the Pleiades Lustre filesystems /nobackup and the Columbia CXFS filesystems /nobackup are mounted on the bridge nodes. To copy files between the Pleiades Lustre and Columbia CXFS filesystems, log into a bridge node and use the `cp`, `mcp`, or `shifto` command to perform the transfer.

## File Transfers to Mass Storage

The Pleiades /nobackup filesystems are mounted on Lou2. Thus, the easiest way to transfer files between Pleiades and Lou2 is to initiate a command such as **cp**, **mcp**, **tar**, or **shifc** on Lou2. For example:

```
lou% mcp /nobackup/username/foo $HOME
```

If you initiate the transfer on Pleiades, the commands **scp**, **bbftp**, **bbscp**, and **shifc** are available to do the transfers between a Pleiades front-end or bridge node and Lou. Since **bbscp** uses almost the same syntax as **scp**, but performs faster than **scp**, we recommend using **bbscp** in cases where you do not require the data to be encrypted. For very large file transfers, we recommend the Shift utility, developed at NAS.

See also [File Transfer Overview](#), and [File Transfer Commands](#).

File transfers from the Pleiades compute nodes to Lou must go through one of the PFEs or bridge nodes first. See [Streamlining File Transfers from the Pleiades Compute Nodes to Lou](#) for more information.

When sending data to Lou, keep your largest individual file size under 1 TB, as large files will occupy all of the tape drives, preventing other file restores and backups.

## Additional Restrictions on Front-end Systems

- No MPI (Message Passing Interface) jobs are allowed to run on the PFEs or the bridge nodes
- A job on bridge[1-2] should not use more than 56 GB; when it does, a courtesy email is sent to the owner of the job
- A job on bridge[3-4] should not use more than 192 GB; when it does, a courtesy email is sent to the owner of the job

Before starting a large-memory session, it is a good idea to check to make sure there is enough memory available. You can run the command **top**, hit "M", and check under the "RES" column for other large memory applications that may be running.

---

Article ID: 181

Last updated: 11 Jan, 2013

Policies -> Computing Policies -> Pleiades Front-End Usage Guidelines

<http://www.nas.nasa.gov/hecc/support/kb/entry/181/?ajax=1>