Basic Tasks

Once you have completed preliminary tasks and read the Environment Overview, you can get started with using the supercomputers. The basic tasks below cover a typical user workflow cycle.

Even these basic tasks require numerous setup steps before you can perform specific tasks on the supercomputers. For each task, follow the links provided below to other Knowledge Base articles to get step-by-step instructions and more details on our HPC environment.

Logging in for the First Time

Once you have obtained your default NAS password and enabled your RSA SecurID token, you must complete a series of steps to be able to log into the supercomputers. If you have not yet obtained your default password and enabled your RSA SecurID token, please see Before You Begin.

After you log in for the first time, subsequent logins will be more straightforward. The first time you access a system in the secure enclave, you will log into one of the secure front-end (SFE) systems, designated sfe[*6-9*], as shown in the environment overview diagram. You will complete two-factor authentication using your RSA SecurID token and your NAS password. For step-by-step instructions and examples, see Logging into NAS Systems for the First Time.

Note: The SFEs are used only as gateways to the enclave resources and are not to be used for storing data. You can store data in your Pleiades home directory, your /nobackup directory, or on your Lou directory.

Subsequent Logins

Once you successfully log in for the first time, your subsequent logins can be done using either a two-step connection method or a one-step method, both of which require the use of your RSA SecurID token and connection through the SFEs.

TIP: Use the one-step connection method described below. This method takes more time to set up, but saves time in the long run.

Two-step Connection

Initially, you may want to use the two-step connection method, which involves connecting from your desktop first to an SFE and then to resources within the secure enclave, without setting up the SSH passthrough feature. You can choose one of several options for this approach: one uses your NAS password; one uses your NASA personal identity verification (PIV) badge (if you have one); and one requires that you first set up public key authentication. All three also use your RSA SecurID passcode. For instructions and examples, see the following articles:

- Using NAS Password + RSA SecurID Passcode
- Using Public Key + RSA SecurID Passcode
- Using NASA PIV Badge + RSA SecurID Passcode (NASA badge-holders only)

The two-step connection method can be cumbersome and time-consuming if you use several logins, so we recommend setting up the SSH passthrough feature so you can use one-step connection, as described below.

One-step Connection

With this method, you can bypass logging into an SFE each time by using SSH-tunneling to transparently "pass through" the SFEs. The SSH passthrough feature allows you to log into any system in the enclave by typing just one SSH command. To use one-step connection, you will need to first set up public key authentication and SSH passthrough. Once these steps are completed, you can access any of the systems in what appears to be a single step. For step-by-step instructions and examples, see the following articles:

1. Setting Up Public Key Authentication
2. Setting Up SSH Passthrough
3. One-Step Connection Using Public Key and Passthrough

Transferring Files From Your Local System to NAS Systems

You can choose from several different methods to transfer files from your local system to your directories within the secure enclave. Remote file transfer methods and commands such as **`sup/shiftc`** and **`scp`** are supported on most NAS systems. You may need to install the client and/or server software for the secure unattended proxy (SUP)/Shift client or SCP script on your local host in order to use these methods.

TIP: Use the **`shiftc`** command to transfer files from your local system to NAS systems. This command dynamically selects the highest performing transfer method available, and is normally the optimal tool for transferring files of any size.
When you initiate file transfers from your local system to the NAS systems, you will be prompted for your RSA SecurID passcode. (You may then be prompted for further authentication, depending on whether you have set up SSH passthrough.) File transfers that are initiated from the NAS systems to your local system might be a simpler way to transfer files, if your local system does not require two-factor authentication. For more information, see the following articles:

- Remote File Transfer Commands
- Shift File Transfer Overview

Below are examples of initiating **`shiftc`** and **`scp`** from your local system. The examples assume that you have downloaded the SUP client and that your username for your local system is different than your username for the NAS systems.

Note: If your local host username and your NAS username are the same, you can omit *nas_username@* from the command line.

shiftc example

```
your_local_system% sup -u nas_username shiftc lou.nas.nasa.gov:filename ./filename
```

scp example

```
your_local_system% scp nas_username@lou.nas.nasa.gov:filename ./filename
```

Compiling Your Code

Several compilers and tools for major programming languages are available on the NAS systems. Because there is no default compiler, you must load a compiler module. Use the **`module`**

avail command to view all available modules.

Note: NAS recommends using Intel compilers for building your applications on Pleiades, Aitken, Electra, or Endeavour. Available modules include Intel compilers for both Fortran and C++.

You can use the following commands to invoke the various versions of Intel compilers:

ifort
: Intel Fortran

icc
: Intel C

icpc
: Intel C++

For example, to compile the application *filename.f* with **ifort**, type:

```
% ifort filename.f
```

Read the **ifort**, **icc**, and **icpc** man pages for options to use. Type **man *topic*** to access the man page for any particular topic. See Intel Compiler for more information.

Porting and Running MPI Codes with HPE's MPT Library

HPE's Message Passing Interface (MPI) is a component of the Message Passing Toolkit (MPT), which supports parallel programming across a network of computer systems. To use MPI, you must first load an Intel compiler module and an HPE module on Pleiades, Aitken, Electra, or Endeavour. The example below shows how to compile and link your code with MPT:

```
% ifort -o filename filename.f -lmpi
```

To run your MPI codes with MPT in a PBS job, insert the following in your PBS script, where *xx* represents the number of MPI processes:

```
mpiexec -np xx filename
```

For more information about using MPI and MPT, see HPE MPT. For a comprehensive overview of compiling and porting code, see the articles listed in Porting and Developing Overview.

Submitting and Running Jobs

Once you have completed the basic setup, transferred your files to the NAS systems, and compiled your code, you're ready to run jobs on the compute nodes.

NAS facility supercomputers use the Portable Batch System (PBS) from Altair Engineering, Inc., for job submission, monitoring, and management. You can run a job through an interactive session using PBS, or you can submit batch jobs by using a PBS job script. In either case, all jobs are submitted to PBS by issuing the **qsub** command from the PFEs.

For more information about PBS, see Portable Batch System (PBS) Overview. For more information about all of the PBS commands described in the following sections, see Commonly Used PBS Commands.

Submitting Interactive Jobs

For interactive jobs, use the PBS command **`qsub -I`** (interactive) as shown in this section.

Running Interactive Jobs on Pleiades, Electra, or Aitken

To run a job interactively on Pleiades, Electra, or Aitken compute nodes, use the **`qsub -I`** command as shown in the following example. Note that in the command line, you can request specific resources by using the **`-l`** (limits) option.

```
pfe21 qsub -I -lselect=1:ncpus=24:model=has,walltime=1:00:00
```

In the example, the user submits a job from pfe*21*, requesting one Haswell (**`has`**) node with 24 cores and 1-hour maximum wall-clock time. Model types for all three systems are shown in the following table:

Aitken	AMD Rome	**`rom_ait`**
	Cascade Lake	**`cas_ait`**
Electra	Skylake	**`sky_ele`**
	Broadwell	**`bro_ele`**
Pleiades	Broadwell	**`bro`**
	Haswell	**`has`**
	Ivy Bridge	**`ivy`**
	Sandy Bridge	**`san`**

Running Interactive Jobs on Endeavour

To submit a job to Endeavour, specify an Endeavour PBS queue served by pbspl4, in the command line. For example:

```
pfe21% qsub -I -lselect=2:ncpus=28:mem=185G:model=cas_end -q e_normal@pbspl4
```

Submitting Batch Jobs

Before submitting a batch job, you must first create a PBS job script. The following sample script uses the **`select`** option to request one Pleiades Haswell node with 24 cores per node and 1 hour of maximum wall-clock time:

```
#PBS -l select=1:ncpus=24:model=has
#PBS -l walltime=1:00:00

cd $PBS_O_WORKDIR

mpiexec -np 24 ./a.out
#end of script
```

Once you've created your script, you can submit it to run on the compute nodes. For example, to submit a job from pfe*21* to the Pleiades, Electra, or Aitken compute nodes, issue the following command:

```
pfe21% qsub job_script
```

To submit a job from pfe*21* to Endeavour, specify an Endeavour PBS queue served by pbspl4, in the command line. For example:

```
pfe21% qsub -q e_normal@pbspl4 job_script
```

Monitoring Jobs with myNAS and qstat

You can monitor your jobs via the myNAS web portal or mobile applications, or by using the **`qstat`** command.

Using myNAS to Monitor Jobs

The myNAS portal is a website that enables you to easily monitor your PBS jobs and check allocation and usage for your projects. For more information, see Using the myNAS Portal.

You can also download and run the myNAS mobile applications on your iOS or Android device to view job status, receive notifications when jobs change state or produce output, and view output files. For more information, see the following articles:

- Installing the myNAS Mobile App
- Using the myNAS Mobile App

Using the qstat Command to Monitor Jobs

When you submit a job to PBS, it is assigned a job identification number (job ID). PBS jobs are usually in either the Running (R) state or the Queued (Q) state. Use the **`qstat`** command as follows to find the status of your jobs, using job ID 12345 as an example:

`qstat`
: show the status of your jobs

`qstat -nu` ***`nas_username`***
: list all jobs that belong to you

`qstat 12345`
: check the status of a job

`qstat -f 12345`
: check the status of a job (with full details)

`qstat -s 12345`
: show why a job is not yet running

Managing Jobs

If you need to delete a job or place it on hold, use the following PBS commands, which use job ID 12345 as an example:

`qdel 12345`
: delete (cancel) a job

`qhold 12345`
: place a job on hold

`qrls 12345`
: release a job from hold

For troubleshooting information, see the following articles:

- Common Reasons for Being Unable to Submit Jobs
- Common Reasons Why Jobs Won't Start

Managing Your Data

Managing the data that your jobs use and create is an integral part of computing at NAS. At different times you may need to transfer data to/from NAS to your remote site, check your disk use and quotas, copy files from the Lou tape drives, or share files with other NAS users. For more information, see the following articles:

- Remote File Transfer Commands
- Pleiades Home Filesystem
- Pleiades Lustre Filesystems
- The dmget Command
- Using 'Giveto' to Copy Files and/or Directories to Another User

Checkpointing Jobs

A checkpoint is a snapshot of your job taken periodically while it is running. If the job runs out of wall-clock time or fails for any reason, it can be restarted from the checkpoint. Be aware that none of the compute node operating systems have an automatic checkpoint capability. Therefore, for jobs that need a lot of resources and/or long wall-clock time, we recommend implementing a checkpoint/restart capability in your source codes or job scripts.

In the event that a system crashes or experiences certain system issues, PBS automatically reruns jobs when the system is back up and issues are resolved. If you do not want PBS to rerun your job, add the following line to the top of your PBS script:

```
#PBS -r n
```

Storing Your Files

To ensure your data is stored securely, and to maintain directory space and avoid quota limits, you will periodically transfer your files from the scratch/nobackup directory and the home directory to the Lou mass storage systems.

The scratch /nobackup filesystems are mounted on Lou, enabling disk-to-disk copying, which provides the highest transfer rates. You can use the commands **`shiftc`**, **`cp`**, or **`mcp`** to copy files or make tar files directly from your scratch/nobackup directory to your Lou home directory. Note that Pleiades, Aitken, Electra, and Endeavour share the same home and /nobackup filesystems.

Network file transfers to the Lou systems are sent through the 10 GigE interface by default. You can use the commands **`shiftc`** and **`scp`** to transfer your files. We recommend using **`shiftc`**, as it dynamically selects the highest performing transfer method available.

File transfers from the compute nodes to Lou must first go through the front-end nodes (PFEs for Pleiades, Aitken, Electra, and Endeavour) or through the /nobackup filesystems. For more information, see the following articles:

- Local File Transfer Commands
- The Lou Mass Storage System

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

Post-Processing Your Data

After your work has completed running on the compute nodes, you can perform post-processing tasks on the Lou data analysis nodes.

The NAS hyperwall visualization system provides a supercomputing-scale resource to process very large datasets produced by the compute nodes, and NASA scientific instruments. Although general users cannot access the system, our visualization experts provide post-processing services. For an overview of the system, see Visualization System: hyperwall.

Article ID: 264
Last updated: 28 Apr, 2022
Revision: 116
New User Orientation -> Basic Tasks
https://www.nas.nasa.gov/hecc/support/kb/entry/264/