The NAS New User Orientation guide gives first-time users a brief overview of the high-performance computing (HPC) environment operated by staff in the NASA Advanced Supercomputing (NAS) facility at Ames Research Center, Moffett Field, California. It contains all the basic information you must know to get started with using the supercomputing resources and services provided at the facility. Returning users may find the guide useful as a quick refresher.

Throughout this guide you will find links to other NAS Knowledge Base articles to help you become familiar with some of the details needed to navigate the complex NAS environment. You can also use the Knowledge Base to browse FAQs and search for other topics of interest. The resources within the NAS environment are funded through NASA's High-End Computing Capability Portfolio, and the terms "NAS" and "HECC" are used interchangeably throughout the Knowledge Base.

The guide includes the following basic topics:

- **Introduction**: Covers what you need to know before you begin, some reminders about your responsibilities as a NAS user, and typographic conventions used.
- **Environment Overview**: Describes the general setup of the NAS/HECC secure environment, including the computing, network, and data storage components and the user environment.
- **Basic Tasks**: Provides a general description of the tasks that comprise a typical user workflow.
- **Understanding and Managing Your Allocations**: Summarizes how computing time is calculated, tracked, and charged.
- **Helpful Resources**: Provides links to help you find more information and summarizes additional services that are provided by NAS experts.

If you still can't find the information you're looking for, the NAS Control Room staff are available 24x7 to provide additional help.

Toll-free Telephone: 800-331-8737
Local Telephone: 650-604-4444
E-mail: support@nas.nasa.gov

Before You Begin

To make the best use of this guide and to perform basic tasks, you must already have done the following:

- Received approval for your request for allocation
- Completed the two-phase user accounts process for NAS systems
- Enabled your RSA SecurID token for two-factor authentication

If you need assistance with any of these items, call the NAS Control Room at the telephone number shown above.

User Responsibilities

When you completed your Account Request Form, you agreed to abide by the NAS Systems Environment Rules of Behavior. Below are a few basic reminders.

- The computing systems are unclassified systems; classified information may not be entered, processed, or stored.
- Users cannot divulge access information (for example, lists of user accounts).
- Users cannot share their account(s) with anyone, including sharing account passwords, providing access via a .rhost entry, or any other means of sharing. If you are found sharing accounts, your accounts will be disabled.
- It is your responsibility to become familiar with various NASA security-related policies.

Conventions Used in This Guide

The following table shows the typographical conventions used in this guide.

Convention	Description	Examples
Percent sign (%)	Indicates the system prompt; your system prompt may be different	pfe*21*% Local%
Monotype	Indicates commands	Run the command mylou

Convention	Description	Examples
Italicized words or phrases	Indicate command variables; variable number ranges; substitute your own phrase; rarely, indicates emphasis or introduces a new term	*tbd* lfe[*5-8*] *your_username* *Never* share your password. The *hotspots* analysis ...
Bold	Indicates man pages to reference	Read the **ifort** man page for more information

Our HPC environment is operated by staff in the NASA Advanced Supercomputing (NAS) Division at Ames Research Center located at Moffett Field, CA. The supercomputers and support staff are funded by NASA's High-End Computing Capability (HECC) Project.

The topics in this section summarize the system components, such as the compute nodes, secure network connections, front-end systems, and data storage facilities; and the user environment, including information about the Linux operating system and the various filesystem directories that are available for your use.

Systems Overview

The systems available for your use are all protected within a secure environment. The figure below shows both the physical connections among all the components and their functional relationships through a typical user workflow. The environment is described in more detail in the following sections.

The Secure Enclave

NAS supercomputing components are protected within a secure enclave that can be accessed only by authenticated users through the following secure bastions:

- Secure front ends (SFEs)
- Secure Unattended Proxy (SUP)

Components protected within the secure enclave include:

- Supercomputers: Athena, Aitken, Cabeus, Electra, Endeavour
- Front-end nodes:

  Note: Although the Pleiades compute nodes have been retired, the PFEs continue to serve as the front ends for Electra and some Aitken compute nodes.

  - Athena front ends (ATHFEs), used by the Athena nodes
  - Pleiades front ends (PFEs), used by Electra, Endeavour, and Aitken Cascade Lake and Rome nodes
  - Aitken front ends (AFEs), used by Aitken Milan nodes
  - Cabeus front ends (CFEs) and Cabeus Grace Hopper front ends (CGHFEs), used by Cabeus
  - Lou front ends (LFEs), used by Lou
- Lou mass storage system
- Lou data analysis nodes (LDANs)
- hyperwall visualization system

The following sections give an overview of each component of the secure enclave.

Bastions

Secure Front Ends (SFEs)

The secure front ends (SFEs) provide inbound connection from your local system to the secure enclave. The first time you access the HECC systems within the enclave, you will authenticate through an SFE. Subsequently, you can use any of the bastions to access systems in the enclave.

For an overview of the initial authentication process, see Logging in for the First Time. For more information about the SFEs, see the article Role of the Secure Front Ends.

Secure Unattended Proxy (SUP)

The Secure Unattended Proxy (SUP) allows you to pre-authenticate to the secure enclave for one-week periods, during which you can perform unattended (batch) file transfers. After you complete the setup process, SUP is the most efficient and convenient method for transferring files from your remote system. For more information, see Using the Secure Unattended Proxy (SUP).

Compute Front Ends

The heterogeneous HECC environment uses several different front-end nodes to access the various node types.

The Electra and Endeavour supercomputers—along with Aitken's Cascade Lake and Rome nodes—share the Pleiades front-end systems (PFEs). You can use the front ends to edit files, compile your code, run short debugging and testing sessions, and submit batch jobs to the compute nodes.

Note: For Athena, use the Athena front-end systems (ATHFEs). For the Aitken Milan compute nodes, use the Aitken front-end systems (AFEs). For the Cabeus A100 GPU nodes, use the Cabeus front-end systems (CFEs). For the Cabeus Grace Hopper GH200 nodes, use the CGHFE front- end systems.

See the following articles for more information:

- PFE Usage Guidelines
- Front-End Load Balancer Among the PFEs

Compute Nodes

There are currently five supercomputers available for users: Athena, Aitken, Cabeus, Pleiades, Electra, and Endeavour.

Athena

Athena, NASA's most powerful supercomputer, is housed in Module 2 of the Modular Supercomputing Facility, an environmentally-friendly module located a short distance from the main NAS building. See the following articles for more information:

- Athena Configuration Details
- Athena Turin Processors
- Preparing to Run on Athena Turin Nodes

Aitken

Aitken is housed in Module 1 of the Modular Supercomputing Facility. Aitken uses the Pleiades front ends (PFEs), filesystems, PBS server, and job queues. See the following articles for more information:

- Aitken Resource Page
- Aitken Configuration Details
- Preparing to Run on Aitken Cascade Lake Nodes
- Preparing to Run on Aitken Rome Nodes

Cabeus

Cabeus is NASA's latest GPU system, currently comprised of 128 NVIDIA GPU A100 nodes and 350 NVIVIA GH200 Superchip nodes. Cabeus uses the Cabeus and Cabeus Grace Hopper front-end systems (CFEs and CGHFEs). See the following articles for more information:

- Cabeus GPU Nodes
- Requesting Cabeus GPU Resources

Electra

Electra is NASA's first prototype modular supercomputing system, housed near the main NAS building. Electra uses the Pleiades front ends (PFEs), filesystems, PBS server, and job queues. See the following articles for more information:

- Electra Resource Page
- Electra Configuration Details
- Preparing to Run on Electra Skylake Nodes
- Preparing to Run on Electra Broadwell Nodes

Endeavour

Endeavour is an HPE Superdome Flex system that provides resources for user applications needing access to large cache-coherent, global shared-memory capabilities in a single system image (SSI). Endeavour uses the Pleiades front ends (PFEs) and filesystems, and shares some of the Pleiades InfiniBand fabric. However, Endeavour uses its own designated Portable Batch System (PBS) server and job queues. See the following articles for more information:

- Endeavour Resource Page
- Endeavour Configuration Details
- Preparing to Run on Endeavour

Mass Storage System

The NAS facility provides long-term storage space on a single mass storage system, known as Lou. This HPE/SGI system has 7.6 petabytes (PB) of disk space and is capable of storing up to 1040 PB (1 exabyte) on tape. See the following articles for more information:

- Your Mass Storage Directory
- Lou Mass Storage System

Post-Processing Systems

Systems provided for post-processing include the Lou data analysis nodes (LDANs), designated as ldan[*11-14*], and the hyperwall visualization system. For a summary on using these systems for post-processing work, see Post-Processing Your Data. For detailed information, see also the following articles:

- PFE Usage Guidelines

- Lou Data Analysis Nodes
- Visualization System: hyperwall

Networks

The NAS high-speed network (NASLAN) includes a 10 gigabit-per-second (Gb/s) local area network and 10 Gb/s peering with other high-speed networks such as the NASA Integrated Communications Services (NICS), Internet2, and the Consortium for Educational Networks in California (CENIC). For an overview, see Networking Resources.

To access the HECC resources inside the secure enclave, you will use the SSH protocol to connect from your desktop system to a bastion (usually the SFEs) through a wide area network and the NASLAN.

User Environment

All HECC systems run the Linux operating system. If you are new to Linux, you can find a lot of helpful information at the user-supported community website Linux.org, including a Beginners Learning Course that provides instruction on the basic directory structure of Linux, how to get around in the directories, how to access Linux documentation (man pages), useful commands, and much more. You can also find support at the Linux forum.

When your NAS account is created, your default Linux shell is set to be the C shell (csh); this is assumed to be the case throughout this guide. If you want to use a different shell as your default, such as bash, call the NAS Control Room staff at (800) 331-8737 or (650) 604-4444 or send an email message to support@nas.nasa.gov to request the change. After the change is made, the new default shell of your choice applies to all of your jobs.

Once you complete the initial setup for your NAS account, you will have access to all the front-end systems described above (PFE, AFE, CFE, CHGFE), the NAS compute resources home filesystems, the Lou mass storage filesystems, and the scratch (/nobackup) filesystems. Your NAS account is authorized to run jobs on all HECC CPU-only compute systems if your project has been granted the HECC Standard Billing Units (SBU) allocation. To run jobs on the HECC GPU compute nodes, your project needs to have been granted a separate GPU SBU allocation. To see if you have the HECC and/or GPU allocations, run the command `acct_ytd` on any NAS front-end system.

NAS supercomputers use the Portable Batch System (PBS) from Altair Engineering, Inc., for job submission, monitoring, and management. For more information about PBS, see Submitting and Running Jobs.

Filesystems

Athena, Aitken, Cabeus, Electra, and Endeavour share the same home and scratch (/nobackup) filesystems. When you log into any of the NAS compute resources front-end systems (such as a PFE, ATHFE, AFE, CFE, CGHFE), you will have access to the following directories:

- A home directory on the NAS compute resources home filesystem, which you can use to store a small number of files such as source code, small input files, and so on
- A /nobackup directory on a Lustre filesystem, which you can use to temporarily store larger files for reading and writing large amounts of data while running jobs

For long-term data storage, you also have access to a home directory on the Lou mass storage systems. The /nobackup filesystems are mounted on Lou, so you can easily copy files there directly from your /nobackup directory.

The NAS compute resources home filesystem and Lou home filesystems are backed up each night. These backups are stored for approximately one year. The scratch (/nobackup) filesystems are *not* backed up.

Quota limits are enforced on all filesystems. Two kinds of quotas are supported:

- Limits on the total disk space occupied by your files
- Limits on the number of files you can store, irrespective of size; for quota purposes, directories count as files

See Quota Policy on Disk Space and Files for more information.

Your NAS Compute Resources Home Directory

Your NAS compute resources home directory is accessible from Aitken, Cabeus, Electra, and Endeavour. Use your home directory to store a limited number of smaller files such as source code and input files. For temporary, short-term storage of larger files, use your /nobackup directory. For long-term data storage, use the Lou mass storage systems. See Home Filesystem for NAS Compute Resources for more information.

Your Scratch (/nobackup) Directory

Use your /nobackup directory to temporarily store large files that read and write large amounts of data when you run jobs. Your /nobackup directory resides on one of several Lustre filesystems, designated /nobackupp*X*. To find out which Lustre filesystem your /nobackup directory is located on, run:

```
pfe21% ls -ld /nobackup/your_nas_username
```

The /nobackup filesystems are also mounted on the Lou mass storage system, so you can access data in your /nobackup directory from Lou without going through Aitken, Cabeus, Electra, or Endeavour.

WARNING: The /nobackup filesystems mean just that: they are not backed up. While this is stating the obvious, some users have lost important data by storing it on these systems over long periods of time. It is your responsibility to copy essential data to either your home directory, to archival storage on the Lou systems, or to your remote system.

See NAS Lustre Filesystems for more information.

Your Mass Storage Directory

For safe, long-term data storage, transfer your files to your Lou home directory. The Lou mass storage filesystem allows you to retrieve your stored files quickly and securely whenever you need them.

You can log into the Lou system just as you would any other HECC system and save data to mass storage by copying your files to your Lou home directory: Lou:/u/*your_nas_username*.

There is no specified data-size quota for your Lou home directory, but you can store up to 250,000 files.

For more information, see The Lou Mass Storage System.

The Data Migration Facility

Data stored on Lou is migrated to tape, as needed, to make space on the disks for more data. Migrated files are retrieved to active disk when you attempt to read or write to them. These migration and retrieval processes are managed by HPE/SGI's Data Migration Facility (DMF), which also enables you to manually list, put, find, and get files that are on tape.

When your data is migrated to tape, two copies are written to two separate tape media in automated tape libraries located in two different buildings. See Data Migration Facility Commands for more information.

The Lou Filesystem

Lou is composed of the Lou front ends (LFEs), designated lfe[*5-8*]. The /nobackup filesystems are mounted on the LFEs, so you can easily copy files there directly from your /nobackup directory.

Although you cannot perform post-processing tasks on the LFEs, the Lou data analysis nodes provide PBS resources to perform post-processing tasks on your Lou mass storage data. For more information, see the following articles:

- The Lou Mass Storage System
- Lou Data Analysis Nodes

Software Modules and Packages

HECC provides many software programs such as compilers, pre- and post-processing programs, analysis tools, and math and scientific libraries. The software is managed through the use of packages and modules that you can load into your home directories.

For more information, see the following articles:

- Software on NAS Systems
- Customizing Your UNIX Environment

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[*6-8*], 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 NAS compute resources home directory, your Lustre /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 CPU-only applications for running on CPU-based nodes on Aitken, Electra, or Endeavour. Available modules include Intel Compiler Classic for C (icc), C++ (icpc), Fortran (ifort) and LLVM-based Intel OneAPI Data Parallel C++ (DPC++)/C++ Compiler (icx for C, icpx for C++) and Intel Fortran Compiler (ifx).

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

ifort or ifx
: Intel Fortran

icc or icx
: Intel C

icpc or icpx
: 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.

Note: Intel Classic C/C++ Compilers (icc, icpc) have been discontinued since Intel 2024 releases. The Intel Classic Fortran Compiler (ifort) has been discontinued since Intel 2025 releases. They are replaced with Intel oneAPI Compilers (ifx, icx, and icpx).

Note: To build GPU applications for running on GPU nodes of Cabeus, we recommend using the NVIDIA compilers, including the NVIDIA HPC compilers (nvfortran, nvc, and nvc++) and the NVIDIA CUDA C/C++ Compiler - nvcc. For more information, see Programming for NVIDIA GPUs.

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 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 an appropriate front-end system:

- PFE for Endeavour, Electra bro-ele, sky-ele, and Aitken cas_ait and rom_ait nodes.
- AFE for Aitken mil_ait nodes.
- CFE for Cabeus mil_a100a nodes
- CGHFE for Cabeus gh200 nodes

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 Electra or Aitken

Note: Jobs submitted for running on Electra and Aitken are served by the PBS server pbspl1, which is set as the default server on the PFEs or AFEs.

To run a job interactively on Electra or Aitken compute nodes, use theqsub -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=40:model=sky_ele,walltime=1:00:00
```

In the example, the user submits a job from pfe*21*, requesting one Skylake (sky_ele) node with 40 cores in the Electra cluster and 1-hour maximum wall-clock time. Model types for all three systems are shown in the following table:

Note: PBS jobs for running on Aitken AMD Milan processors can only be submitted from the front-ends afe01 and afe02, not the PFEs.

Aitken	AMD Rome	rom_ait
	Cascade Lake	cas_ait
	Milan	mil_ait
Electra	Skylake	sky_ele
	Broadwell	bro_ele

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
```

Running Interactive Jobs on Cabeus

Note: : Running jobs on the NAS GPU nodes requires a GPU SBU allocation. See Requesting Cabeus GPU resources.

All Cabeus GPU nodes (except the Grace-Hopper GPU nodes), including the V100 GPU-based sky_gpu, cas_gpu compute nodes, and the A100 GPU-based mil_a100 compute nodes are served by the PBS server pbspl4. These various GPU nodes are accessible via different PBS queues. See Requesting GPU Resources. The example below shows how to request an interactive session with exclusive access to a mil_a100 node using the gpu_devel queue:

```
pfe21% ssh pbspl4
pbspl4% qsub -I -q gpu_devel -lselect=1:ncpus=64:ngpus=4:mem=500GB:model=mil_a100 -lplace=scatter:excl
```

The newly deployed Cabues Grace-Hopper GPU nodes are served by the PBS server pbs05a.

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 Electra Skylake node with 40 cores per node and 1 hour of maximum wall-clock time:

```
#PBS -l select=1:ncpus=40:model=sky_ele
#PBS -l walltime=1:00:00

cd $PBS_O_WORKDIR

mpiexec -np 40 ./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 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 by using theqstat 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.

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
- NAS Compute Resources Home Filesystem
- NAS 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 Aitken, Electra, Endeavour, and Cabeus 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 commandsshiftc 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 Aitken (Cascade and Rome nodes), Electra, and Endeavour
- AFEs for Aikten Milan nodes

- CFEs for for the Cabeus mil_a100 GPU nodes
- CGHFEs for Cabeus gh200 GPU nodes

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 theLou 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.

As part of applying for a NAS account, you (or your principal investigator) requested an allocation for computing time on HECC systems. Allocations for NAS supercomputing resources are granted per project group ID (GID), and are specified in standard billing units (SBUs).

Usage Charge Methods

System usage is charged as follows:

- **Front-End Systems**: Usage is *not* charged on the front-end systems (such as the PFEs, AFEs, CFEs, CGHFEs, and LFEs). However, remember that these front-end nodes are intended for editing and/or compiling and running short testing jobs—not for running production jobs. If you misuse these systems, your jobs will be terminated.
- **Lou Data Analysis Nodes**: Usage is *not* charged on the Lou data analysis nodes (LDANs), which provide dedicated PBS resources for pre- and post-processing tasks.
- **Compute Nodes**: Usage on the Aitken, Cabeus, Electra, and Endeavour compute nodes*is* charged. The number of SBUs charged to a job is calculated by multiplying the number of total wall-clock hours used by the minimum allocation units (MAUs)—the smallest unit of hardware resource that the Portable Batch System (PBS) will allocate to a job.

SBU rates vary by processor type and system. See Job Accounting for a full list.

To continue running jobs after your SBU allocation has been expended, you will need to request more hours. You can check your remaining SBU balance by running the acct_ytd command. See Job Accounting Utilities for more information.

TIP: You can easily keep track of allocation and usage for all of your GIDs (and monitor all of your PBS jobs) in one central location by logging into the myNAS Portal.

Charging to a Non-Default GID

If you have access to more than one GID, only one of those GIDs is set as your default. The default GID is listed in the/etc/passwd file of each system you have access to. Your jobs will be charged to the default GID unless you specify a different one. If you want to charge your usage to a non-default GID (for example, s0901), add the GID to your PBS script. For example:

```
#PBS -W group_list=s0901
```

Understanding Queues

All HECC supercomputers use the Portable Batch System (PBS) to manage both interactive and batch jobs. Aitken and Electra use the same PBS server (pbspl1) and job submission queues. Endeavour and Cabeus' A100 nodes have a separate PBS server (pbspl4) and separate queues. Cabeus' GH200 nodes use a third server (pbs05a) and separate queues.

Although different queues are available on different systems, queues typically have constraints on maximum wall-clock time and/or the number of nodes allowed for a job. Some queues have other constraints or are restricted to serving certain users or GIDs. In addition, mission directorate limits are set on the number of cores available on Aitken and Electra to ensure that each mission directorate can access a fair share of resources.

You can use the qstat command as follows to view various types of information:

qstat -Q
: Lists available queues and their constraints on a system.

qstat -Qf
: Lists available queues and their constraints on a system.

qstat -W shares=-
: Lists mission shares available for each mission.

For more information, see the following articles:

- PBS Job Queue Structure
- Mission Shares Policy

See the following pages for important user announcements, system status, and training sessions:

- User News
- System Status
- User Training Webinars
- Past Webinars Archive
- Frequently Asked Questions
- NASA Advanced Supercomputing (NAS) Division

Remember that NAS Control Room staff are available 24 hours a day, seven days a week:

Toll-free Telephone: 800-331-8737
Local Telephone: 650-604-4444
E-mail: support@nas.nasa.gov

The following support services are also offered free of charge to all users:

- Application performance optimization and code porting
- End-to-end network services for troubleshooting and performance
- Customized training and support to help users efficiently manage large amounts of data
- Advanced scientific visualization and analysis capabilities

For more information about our service offerings, see HECC Services Overview.