The PFEs are the front-end systems for Aitken, Electra, and Endeavour. They provide an environment that enables quick turnaround for tasks such as file editing, file transferring, compiling, and short debugging/testing sessions, as well as for batch job submissions via PBS to Aitken, Electra, to Endeavour.

WARNING: The PFEs use Intel Sandy Bridge processors. If you use a PGI compiler to build your executable on the PFEs, be aware that by default the executable is optimized for the Sandy Bridge microarchitecture (which includes Sandy Bridge and Ivy Bridge) and will not necessarily execute on the compute node processors. See PGI Compilers and Tools for information on generating a single executable that will work on all processor types.

You cannot use SSH to access the compute nodes except for the subset of nodes your PBS job is running on.

## Restrictions on PFEs

Default usage quotas are established on the PFEs to support the availability and reliability of the these systems for all users. If you exceed the default usage limits, the memory and processor resources available to you on the PFE will be reduced for a period of time, and you will be notified via email.

The hard quota defaults for each user are 4 cores for processor usage and 8 gigabytes (GB) for memory usage. The soft quota defaults for each user are 2 cores for CPU usage and 4 GB of memory usage.

For details, see Usage Quotas on PFEs.

WARNING: MPI (Message Passing Interface) jobs are *not* permitted to run on the PFEs.

## Pre-Processing and Post-Processing Data

For pre- and post-processing applications such as Tecplot, IDL, and MATLAB, we recommend using a reserved front end (RFE) or the Lou data analysis nodes (LDANs).

You can reserve an RFE by running the "reservable front end" command (pbs_rfe), specifying what type of node you want and for how long. The command makes a PBS advanced reservation for a node and tells you which node is reserved.

You can request the LDANs in the qsub command line. The Lou home filesystems, NAS compute resources home filesystems (/homeX/username), and /nobackup filesystems are all mounted on the LDANs.

Note: There are only four LDAN nodes, so there are often wait times for using them.

## File Transfers to Mass Storage

The /nobackup filesystems are mounted on Lou, so the easiest way to transfer files between NAS compute systems and Lou is to initiate a command such as shiftc, cp, mcp, or tar on Lou. For example:

```
lou% shiftc /nobackup/username/filename $HOME
```

If you initiate the transfer from a PFE, you can use the commands scp or shiftc to transfer files between a PFE and Lou. For very large file transfers, we recommend the NAS-developed Shift tool (shiftc).

File transfers from the compute nodes to Lou must first go through one of the PFEs.

When sending data to Lou, keep your largest individual file size under 1 terabyte (TB). Files larger than 1 TB will occupy all of the tape drives, preventing other file restores and backups.

The secure front ends (SFEs) are the bastion hosts protecting the NAS secure computing enclave. The enclave includes all of the major HECC systems, including: Aitken, Electra, Endeavour, Lou, and the hyperwall.

To access resources inside the enclave from your local system, you must first use Secure Shell (SSH) to connect to one of the SFEs (sfe[*6-8*]) and log in using one of the following two-factor authentication methods:

- RSA SecurID + NAS password
- RSA SecurID + NASA personal identity verification (PIV)
- RSA SecurID + public key

Once authenticated, you can then use SSH to access any of the NAS systems from the SFE.

You can avoid having to log in twice (first to an SFE and then to a system inside the enclave) by setting up SSH Passthrough.

Note: Using SSH to connect from the SFEs to hosts outside of the enclave is not allowed.

## New SFE Features and Changes

The SFEs (sfe[*6-8*]), which replaced some old SFEs in 2021, have some new features as well as some changed functionality:

- The new SFEs are built on a different and more restrictive architecture than the old ones. You can no longer directly read or write any file on the SFEs. You can only execute the commands ssh, ssh-balance, and a new command called ssh-key-init.
- SSH public keys used for SSH passthrough are managed using the new ssh-key-init command. For instructions on using this command to initialize public keys, see: Setting Up Public Key Authentication
- Transfers to and from the SFEs themselves are no longer allowed. Existing mechanisms using SSH passthrough and SUP/Shift are still supported as previously. For transfer instructions, see: Inbound File Transfer Through SFEs: Examples
- When you log into an SFE using the RSA SecurID passcode + NAS password method, the password will now be prompted *before* the RSA SecurID passcode (instead of after).
- A new authentication method, RSA SecurID passcode + NASA PIV, is now available for users who hold NASA personal identity verification (PIV) badges. When your NASA badge is attached to your workstation via a card reader, this method enables you to log in using your PIV personal identification number (PIN) and RSA SecurID passcode.

The PFEs enable quick turnaround for tasks such as file editing, file transfer, compiling, and short debugging/testing sessions, as well as for PBS batch job submissions via PBS to the Aitken, Electra, or Endeavour compute nodes.

To support the availability and reliability of the PFEs for all users, resource-intensive tasks should be submitted to the compute nodes, not run on the PFEs. Therefore, the PFEs have the following default usage quotas:

- Hard quota limits for each user are 4 cores for CPU usage and 8 gigabytes (GB) for memory usage.
- Soft quota limits for each user are 2 cores for CPU usage and 4 GB of memory usage.

If your processes running on a PFE exceed the default usage limits, the memory and CPU resources available to you on the PFE will be reduced for a period of time, and you will be notified via email.

## Email Notifications and Consequences

If you exceed the soft quota usage limits on a PFE for more than 30 minutes, the memory and CPU resources available to you on that PFE will be reduced by 20% for a 30-minute period. During this time, your CPU processes will be throttled, and processes exceeding the memory limit will be terminated.

Resource availability will return to normal after 30 minutes. However, if you exceed the soft quota limits again within a three-hour window, the memory and CPU resources available to you on the PFE will be reduced further—by 50% this time—and the three-hour timer will reset. If you exceed the soft quota limits a third time, the resources available to you on the PFE will be reduced by 70%.

Each time this happens, you will receive an email notification that provides usage details for all of your processes running on the PFE. To stay within the quota limits, consider stopping any resource-intensive processes listed in the email; these processes should instead be run on a Pleiades, Aitken, or Electra compute node via a PBS job submission.

## Whitelisted Processes

Processes such as compilation and file transfers are whitelisted and will not count toward CPU quota limits. However, all processes are still counted toward memory quota limits.