

2.0 Environment Overview

Category: New User Orientation

Columbia Phase Out:

As of Feb. 8, 2013, the Columbia21 node has been taken offline as part of the Columbia phase out process. Columbia22-24 are still available. If your script requires a specific node, please make the appropriate changes in order to ensure the success of your job.

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

The sections below summarize the overall supercomputing environment, including secure network connections, front-end systems, and data storage facilities.

2.1 Capabilities and Components

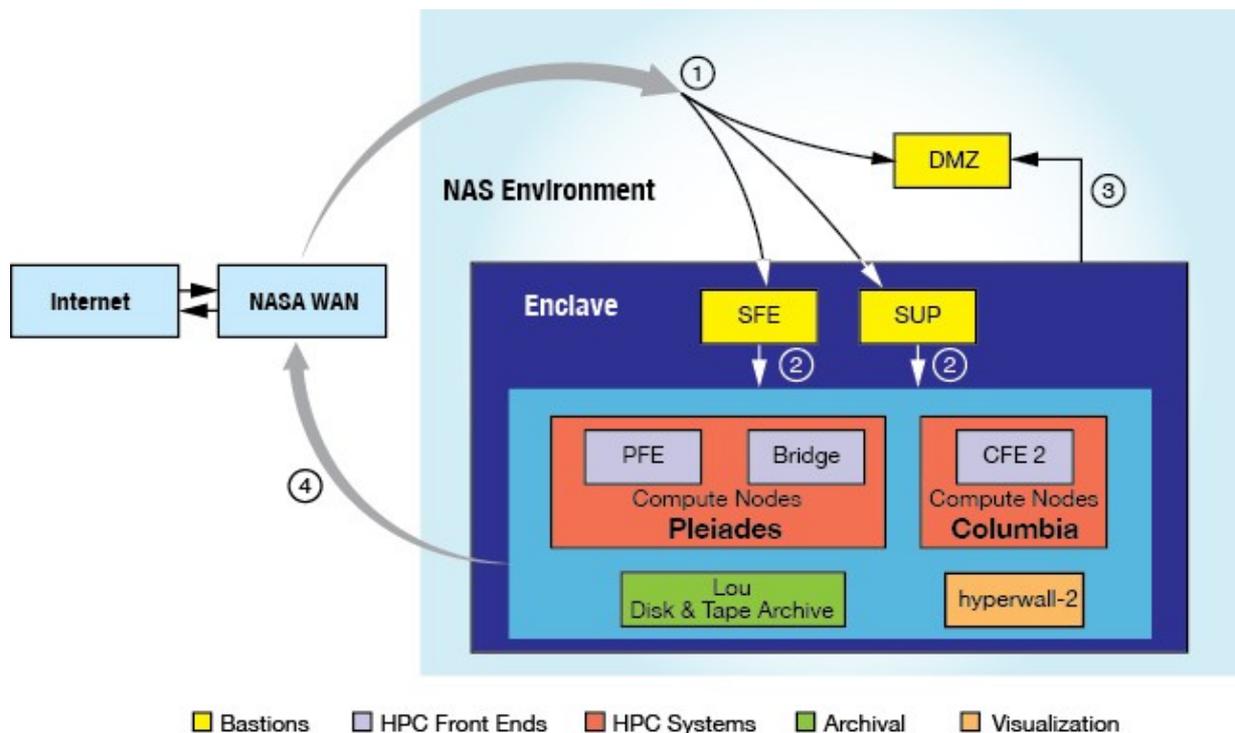
The HPC environment provides the following capabilities and components, including integrated supercomputing support services, augmented by customizable support throughout the entire life cycle of your projects. If you need assistance at any point, please contact the NAS Control Room 24x7: (800) 331-8737 or (650) 604-4444, support@nas.nasa.gov.

2.1.1 Secure Enclave

The enclave (see figure below) is a secure, networked area containing the following:

- **Bastions**
 - ◆ SFEs: Secure Front-End systems, the outermost front ends, provide inbound connection. Two-factor authentication using your RSA SecurID fob is required when issuing commands such as `ssh`, `scp`, and `bbftp` for inbound connection
 - ◆ SUP: Secure Unattended Proxy allows you to pre-authenticate for unattended (batch) file transfers
 - ◆ DMZ: Used for file staging to or from the HPCs
- **High-performance computers (HPCs)**

- ◆ Pleiades supercomputer: Access the Pleiades compute nodes through HPC front ends (the PFEs or bridge nodes. Use the HPC front ends to edit files, compile, run short debugging and testing sessions, and submit batch jobs to the compute nodes. Bridge nodes, with more memory and better interconnects than the PFEs, can also be used for pre- and/or post-processing, and file transfers between Pleiades and Columbia or Lou.
- ◆ Columbia supercomputer: Access Columbia compute nodes through cfe2 or "ssh" directly into compute nodes c21, c22, c23, or c24.
- **Mass Storage Systems:** Lou1 and Lou2, disk and tape archive.
- **Visualization System:** hyperwall-2 provides resources for processing and visualization of large scientific datasets.



NAS environment, components, and workflow:

1. Inbound traffic going to the NAS bastions: the SFEs, SUP, or DMZ
2. Inbound traffic from an SFE or SUP to a Pleiades or Columbia front-end system (bridge nodes, PFE or cfe2)
3. Outbound file transfers through the DMZ
4. Outbound traffic from Pleiades or Columbia directly to remote systems

2.1.2 Supercomputers

The high-end computers at the NAS facility currently include Pleiades and Columbia. While the facility offers a large compute capability, keep in mind that you have been given access

to and allocations on specific systems; not all supercomputing resources are available for your use.

Pleiades

NASA's flagship supercomputer, and one of the most powerful production systems in the world, Pleiades is an SGI ICE cluster containing multiple generations of Intel processors. Bridge nodes with larger amounts of memory and higher-speed network connections compared to the PFEs, offer faster file transfers between Pleiades and Columbia or the Lou mass storage systems.

- Overview of the [Pleiades supercomputer](#)
- Details on the [Pleiades configuration and usage guidelines](#)

Columbia

The Columbia system, the predecessor to Pleiades, currently comprises SGI Altix 4700 compute servers.

- Overview of the [Columbia supercomputer](#)
- Details on the [Columbia configuration and usage guidelines](#)

2.1.3 Visualization System: hyperwall-2

The hyperwall-2 provides users with a supercomputing-scale resource to process very large datasets produced by the HPCs and NASA scientific instruments. Access to this system by general users is limited and requires special authorization.

- Get an overview of the [hyperwall-2](#)

2.1.4 Networks

To access the HPC resources, use [the SSH protocol](#) to connect from your desktop system to a NAS bastion (usually the SFEs) through a wide area network and the NAS local area network.

For remote users, two [DMZ servers](#) are available for staging files between the enclave and the outside world. The DMZ file servers provide limited storage (2.8 TB each) for temporary file storage for very short durations. Files older than 24 hours are automatically removed.

See [Network Tools](#) and [File Transfer Overview](#) for more information.

2.1.5 Mass Storage

The NAS high-performance computing environment includes mass storage systems Lou1 and Lou2, capable of storing up to 115 petabytes (PB) on tape. The HPC compute systems and the mass storage systems combined have over 10 PB of disk storage on the floor.

2.1.6 Other Capabilities and Services

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

Find out more about our service offerings.

NAS User Support staff are available 24x7.

Toll-free Telephone: (800) 331-8737

Local Telephone: (650) 604-4444

E-mail: support@nas.nasa.gov

Article ID: 25

Last updated: 14 Feb, 2013

New User Orientation -> 2.0 Environment Overview

<http://www.nas.nasa.gov/hecc/support/kb/entry/25/?ajax=1>