

High Performance Computing (HPC) and Research Computing Pilot Program

Frequently Asked Questions

Contents

1. What is the HPC and Research Computing Pilot Program:	1
2. How do I request HPC Resources?.....	1
3. What resources are available in this offering?	1
HPC Partner Program through NCSU:.....	2
1. How do I get access to NCSU Partner Program Resources?.....	2
2. How do I connect to the Hazel Cluster?	2
3. How much storage is available?	3
4. How do I transfer files?.....	3
5. How do I submit jobs?	3
6. What compilers are available?	4
7. What software is available on the cluster?	4
8. How do I install software on the cluster?.....	4
9. How do I get support?	4
10. What training, tutorials or other resources are available?	4

1. What is the HPC and Research Computing Pilot Program:

- REDE, in coordination with ITCS, has invested in two pilot programs to provide access to high performance computing (HPC) resources for primary investigators and their labs to generate the preliminary data needed to make high quality applications to extramural grant funding agencies. Initially the Pilot focused on HPC request, but recognizing the broad interest the program was expanded to include research computing.
- Additional information about the Pilot may be found on the REDE website. [High Performance Computing \(HPC\) Program | Research | ECU](#)

2. How do I request HPC Resources?

- To request access please complete the [High Performance Computing Application Form](#)
- Your request will be reviewed monthly, and you will be notified of the status.

3. What resources are available in this offering?

- Currently there are (2) resource offerings. 1.) Microsoft Azure, 2.) NC State HPC Partner Program
 - Microsoft Azure Cloud Resources:
 - Azure's HPC On-Demand Platform
 - Scalable compute resources:

- H-Series systems utilizing AMD EPYC 7003-series CPU cores or Intel Xeon Platinum 8168 CPU cores.
- M-Series memory optimized systems for heavy in-memory workloads and Intel Xeon Platinum 8180M (Skylake) CPU cores
- N-Series systems optimized for GPU workloads including NVIDIA Tesla P40, M60, or V100 GPUs
- Native cloud applications for AI, ML, IoT, etc.
- Slurm job scheduling and queuing, InfiniBand interconnects for clusters, and large storage area for data.
- Azure resources are consumption based and available through June 30, 2024. Resources may be extended depending on funding availability.
- More details on Azure’s HPC capabilities in the cloud can be found here: <https://azure.microsoft.com/en-us/solutions/high-performance-computing>
- Dedicated Nodes through the NCSU HPC Partner Program:
 - REDE has purchased a 3-node cluster totaling 192 CPU cores and 1,280 GB of RAM in NC State’s HPC core facility. The cluster is known as the “Hazel Cluster”
 - The cluster includes a large storage area for data, InfiniBand cluster interconnects, and 25 gigabit Ethernet for network connectivity.
 - IBM’s Spectrum Load Sharing Facility (LSF) is the job queuing and scheduling system.
 - In addition to the 3 compute nodes, ECU has purchased up to 1 million compute hours of time on the Hazel cluster providing access to additional cores and compute queues.
 - Partner receives first priority access to the resources they purchased. When resources are not being used they will be available for use by other HPC cluster users.
 - Shared file system includes a 365TB parallel file system, 95TB of NFS mounted scratch space and a hierarchically managed file system for long term data storage with more than 46TB of disk space and more than 500TB of tape storage.
 - Servers will have access to the shared file system and applications on the cluster. Intel and Portland Group C/C++, Fortran77, and Fortran 90/95/2003 compilers and intel Math Library have been licensed and available to partner servers.
 - Current hardware available is Lenovo servers. OIT maintains the hardware. Typical availability is 99%.
 - NCSU Partner Program resources are available at least till May 2028.
 - More details on NC State’s HPC core facility can be found here: <https://hpc.ncsu.edu/main.php>

HPC Partner Program through NCSU:

1. How do I get access to NCSU Partner Program Resources?

- Once your HPC resource request has been approved, you will be given access to the cluster.
- A new project space will be created, or if requested you can be added to an existing project that you are collaborating.
- A project would consist of the Faculty researcher – PI and their research group.
- To have your research group added to your project, the PI can request this access by emailing the names and PirateID of those needing access to itcs-linux@ecu.edu.

2. How do I connect to the Hazel Cluster?

- Access is provided through a web-based portal named “Open OnDemand”.

- Open a web browser and go to <https://servood.hpc.ncsu.edu/>
- Select “East Carolina University” on the UNC Federation “Select Your Campus” page.
- You will be redirected to the ECU login page allowing you to authenticate your ECU credentials.
- You’ll be presented the “Message of the Day” page along with a menu bar.
 - Hazel Linx Cluster> Files | Jobs | Clusters | Interactive Apps
 - Files - Provides access to your home directory and allows you to upload/download files
 - Jobs – Access to the Job composer and to list active jobs
 - Clusters – Provides shell access
- NCSU has a really good documentation set, so please see [Getting Started on Hazel Cluster](#) for detailed information.

3. How much storage is available?

- **Home Directory:**
 - 1GB per account
 - /home/<login>
 - For housing scripts, small applications, environment files, etc.
 - Daily snapshots and tape backup. (/.snapshot subdirectories)
 - Deleted files are retained for ~ 7 days.
- **Scratch Space:**
 - 20TB per project – parallel file system (GPFS)
 - /share/<project>/<login>
 - For housing data for actively running jobs and analysis
 - Scratch space **is NOT backed up.**
 - Files not accessed in **30 days** are AUTOMATICALLY DELETED.
 - Computation work should be transferred from scratch space promptly to avoid loss of work.
- **Available Storage:**
 - Storage space for application installation will be provided on request.
 - User Apps
 - /usr/local/usrapps/<project>
 - User maintained apps. Need to be available for Linux OS and Intel Xeon Processor
 - The project agrees to maintain the application and address any known vulnerabilities.
 - READ Only from compute nodes. Not intended for data storage
 - Daily snapshots and tape backup. (/.snapshot subdirectories)
- Additional research storage may be purchased. Current cost is \$165 per TB per year and would be funded by your department, grant, or startup funds.

4. How do I transfer files?

- Open OnDemand
 - Simplest method for small files. Limit of ~ 10GB file size.
 - Using the OnDemand web portal, navigate to files and use the menu options.
 - The default location is /home/<user name>. Remember there is a 1GB limit in the home directory so for larger files you’ll need to upload to the /share/<project>/<login> space
 - To change locations, go to Files > Change Directory and type in the desired location.

5. How do I submit jobs?

- Running applications on the login node is not permitted. The user must submit a batch script or request an interactive session via the job scheduler.

- Jobs are submitted from cluster login nodes.
 - Access to login nodes will be through Open on Demand Web portal
 - All Job scheduling must be through LSF – IBM/Platform Computing’s Load Sharing Facility.
 - Queues are available for general use and based on the shared compute hours purchased.
 - Default scheduling priority is fair share and project priority decreases based on number of cores actively in use and recently accumulated project time.
 - For information on creating, submitting and monitoring jobs please refer to [HPC - Running Jobs on the HPC](#)
 - Information on Job Priority may be found at [HPC - Job Priority](#)
6. What compilers are available?
- Three sets of compilers are available on Hazel: GNU, Intel, and Nvidia (formerly Portland Group). Each set includes Fortran, C, and C++ compilers.
 - For more detailed info, please refer to [HPC - Compiling Code](#)
7. What software is available on the cluster?
- Please see [High Performance Computing | Software \(ncsu.edu\)](#) for officially supported applications and libraries
8. How do I install software on the cluster?
- Before installing
 - Review the current install list or see if the module is installed by typing *module avail*
 - Request space in the /usr/local/usrapps/<project> by emailing oit_hpc@help.ncsu.edu using the specific language described on the [HPC user software page](#)
 - Remember custom installed apps must be actively maintained/patched by the user / project group.
 - If you want to use the Conda package management system to assist you in maintaining your software. For information reference [Using Conda](#)
 - For additional guidance on installing software go to [High Performance Computing | Installing software \(ncsu.edu\)](#)
9. How do I get support?
- For general HPC assistance
 - Please open a Team Dynamix ticket for the ITCS Enterprise Linux Team. [Service - Unix/Linux Server Support \(teamdynamix.com\)](#).
 - For NCSU Hazel Cluster support:
 - NCSU Service Now support tickets can be opened by sending email to oit_hpc@help.ncsu.edu.
 - General (non-critical) incidents typically receive a response in two business days.
 - All incidents are addressed according to effort required, with critical incidents receiving priority.
 - To ask a question or for more information refer to [High Performance Computing | Contact Us \(ncsu.edu\)](#)
10. What training, tutorials or other resources are available?
- To learn more about HPC concepts and skills, please see [High Performance Computing | Links for Learning HPC \(ncsu.edu\)](#)
 - An HPC quick start tutorial may be found at <https://hpc.ncsu.edu/guide/>
 - ECU HPC Tutorial (for Python and Conda) - [ECU HPC Tutorial \(for Python and Conda\) - YouTube](#)