

# Yale Center for Research Computing

## HPC Policies

*Effective July 1, 2024*

### Table of Contents

Introduction	1
Access and Accounts	1
Systems Administration	3
Compute & Storage Resources	3
Data & Security Considerations	3
Backups	4
Resource Scheduling and Jobs	5
Research Support and User Assistance	5
Hardware Acquisition and Lifecycle	6

### Introduction

The Yale Center for Research Computing (YCRC) is a computational core facility under the Office of the Provost created to support the advanced computing needs of the research community. The YCRC is staffed by a team of research scientists, computational research support analysts, and systems administrators with expertise in supporting high performance computing (HPC) and computationally dependent disciplines.

The YCRC advances research at Yale by administering a state-of-the-art cyberinfrastructure, providing sustainable research services, and facilitating an interdisciplinary approach to the development and application of advanced computing and data processing technology throughout the research community.

### Access and Accounts

#### Research Access and Accounts

The YCRC operates high performance computing (HPC) resources supporting research in a wide range of disciplines across the university. Upon request, the YCRC will provide a principal investigator (PI) group account on the YCRC's HPC resources to any member of the ladder faculty appointed at the level of Assistant Professor or above, or any member of the research faculty appointed at the level of Research Scientist/Scholar or above who is responsible for their own independent research program. Other university faculty or staff may apply for such an account, subject to additional administrative approval. Once a group account has been created, and with the approval of the PI, members of a PI's research group may request individual user accounts under the group account. The PI is ultimately

responsible for all use of YCRC resources by members of the PI's group, including any costs for data storage and computing that may be associated with such use.

### **Account Deactivation and Removal**

The YCRC audits cluster accounts annually on November 1, and all accounts not associated with a valid NetID will expire at that time. In addition, YCRC needs to be able to contact all account holders for security and communication purposes. Therefore, logins will be disabled throughout the year for accounts that are found not to be associated with a valid email address.

PIs are responsible for notifying the YCRC when a user has left the PI's research group or the University, at which time the user's account will expire unless other arrangements are made with the YCRC. When accounts expire, PIs are responsible for ensuring all files are properly managed or removed. However, the YCRC reserves the right to delete expired accounts and all files associated with them, if necessary. Users should arrange to transfer file ownership when possible before leaving the group(s) with which they have been working.

### **Undergraduate Student Access**

The YCRC focuses primarily on computational research conducted by faculty, research staff, postdocs, and graduate students. However, the YCRC may provide limited access to undergraduate students, most often when they are part of a PI's research group or enrolled in a course that uses YCRC resources. In appropriate circumstances, undergraduate students may request use of YCRC resources for independent research under the oversight of a faculty sponsor/advisor. Such requests will be considered on a case-by-case basis, and approval will be at the discretion of the YCRC.

### **Academic Courses on YCRC Clusters**

Instructors of Yale academic courses may request the use of YCRC facilities for their students. Instructors planning the use of YCRC facilities in their courses are expected to consult with the YCRC and obtain prior approval at least 60 days before the start of the semester. Instructors and their teaching staff will be expected to provide primary support for their students using YCRC facilities, with secondary assistance from YCRC staff as appropriate. The YCRC will do its best to ensure access to its facilities throughout the semester, but scheduled hardware maintenance periods will take place according to the YCRC's published schedule as listed on the [System Status Page](#) on the YCRC website, and instructors should plan accordingly. Instructors should also understand that, in rare instances, clusters may be unavailable due to unforeseen circumstances. All accounts and data created for course use will be disabled and removed 30 days after the end of the semester.

Policies regarding account creation and access to HPC resources are subject to change.

## Systems Administration

The YCRC's engineering group administers all HPC systems, including machines (HPC clusters), storage facilities, and related data center networking. Since these systems are intended to support research applications and environments, they are designed and operated to achieve maximum performance levels and job throughput. The availability of the resources, while important, is secondary to their level of performance. Therefore, the engineering group's primary responsibility is to maximize system stability and uptime while maintaining high performance. To that end, all systems are subject to regular maintenance periods as listed on the [System Status Page](#) on the YCRC website. During these maintenance periods and, rarely, at other times, some or all of the YCRC HPC resources may not be available for use.

## Compute & Storage Resources

The YCRC will provide each PI group access to compute and storage resources on one YCRC cluster, subject to reasonable limits and YCRC discretion. In some circumstances, the YCRC may provide a PI group or some of its members with access to additional resources as appropriate for their computational work. (Such additional access may be temporary.) Users are expected to do their best to use the clusters efficiently and release idle resources.

All users will have access to limited amounts of storage in home, project, and short-term scratch directories, free of charge. Quotas limit the amount of storage and the number of files per user and/or PI group. Users are expected to do their best to delete files they no longer need. Files in short-term scratch directories are purged automatically after 60 days. Using scratch for long-term file storage (through artificial extension of expiration or other means) is forbidden without explicit approval from the YCRC. PIs are responsible for all storage usage by members of their groups. Additional storage may be provided upon request and subject to YCRC discretion, which may or may not incur a cost.

## Data & Security Considerations on YCRC Resources

### User Account Security

User accounts are personal to each user and may not be shared. Under no circumstances may any individual use another person's account. Users are expected to follow standard security practices to ensure the safety and security of their accounts and data.

### Medium or High Risk Data

Users are prohibited from using or storing high or medium risk, sensitive, or regulated data (e.g., HIPAA data, PHI, or PII) at any time on any YCRC facilities other than those explicitly designated for such data.

### Data Use Agreements

Regardless of the sensitivity/risk classification of the data, users are prohibited to use or store, on any YCRC facility, data covered by a data use agreement (DUA) unless the DUA has been approved by the Office of Sponsored Projects and reviewed by YCRC. A copy of the

DUA must be provided to the YCRC and, as early as possible in the process and before storing any covered data on any YCRC facility, the YCRC must be informed of and agree to meet all applicable computing-related requirements of the DUA, including, but not limited to, requirements for data encryption, access control, auditing, and special actions to be taken upon removal of the data. Users who expect to start a project that involves the use of YCRC facilities for sensitive data or data subject to a DUA, or who are applying for external funding for such a project, should consult with YCRC early in the planning process. PIs of such projects are responsible for informing YCRC which users are authorized to access which covered data and for notifying YCRC when there are changes to the list of authorized users.

## **YCRC Security Procedures**

Security of YCRC facilities and all data stored on them is critical, and the YCRC takes several steps to help provide a secure computing environment, including:

- Operating the clusters in a secure data center, with restricted and logged access;
- Using firewalls to allow access only from the Yale campus network or the Yale VPN, and restricting that access to only login and data transfer servers;
- Requiring ssh key pairs (not passwords) for ssh authentication;
- Requiring Multi-factor Authentication for both VPN and HPC authentication;
- Keeping operating systems up to date, and regularly applying security patches;

In addition, users are responsible for the security of YCRC clusters and their data. Accordingly, users are expected to follow standard security practices to ensure the safety and security of their accounts and data. This includes:

- Using strong passphrases on ssh keys;
- Setting permissions appropriately on data files and directories;
- Never sharing private keys, passphrases, or other login information;
- Following the terms of any regulations or Data Use Agreements that cover their data.

Elevated privileges, such as sudo or root access, on all systems operated by the YCRC are strictly limited to YCRC staff to ensure the security and stability of the systems.

Further information regarding Yale cybersecurity and data classifications can be found at <http://cybersecurity.yale.edu>.

## **Backups**

Except as described on the YCRC website for specific clusters, only files in users' home directories are backed up, and then only for a short time (approximately 30 days on most clusters). No other user files are backed up at all. Backups are stored locally, so significant events affecting the HPC data center could destroy both the primary and backup copies of user files. Users should maintain their own copies of critical files at other locations. YCRC cannot guarantee the safety of files stored on HPC resources.

## Resource Scheduling and Jobs

The YCRC's HPC resources are shared by many users. The YCRC uses a workload management system (Slurm) to implement and enforce policies that aim to provide each PI group with fair, but limited, access to the HPC clusters. Users may not run computationally intensive workloads or compilations on the login or transfer nodes. Instead, users must submit such workloads as jobs to Slurm, specifying the amount of resources to be allocated for the jobs. Jobs running for longer than one week are discouraged. Slurm will terminate jobs exceeding their requested resource amounts with little or no warning. To avoid data loss if jobs terminate unexpectedly, users are strongly encouraged to checkpoint running jobs at regular intervals.

Users are expected to abide by the stated purposes and limits of the cluster partitions and submit jobs in alignment with [YCRC best practices](#), such as not running large numbers of very short jobs or workflows that create an excessive number of small files. Jobs found making inappropriate use of a cluster may be canceled without prior notice and repeated offenses after a warning has been issued by YCRC staff can result in account suspension. In extreme cases, where a particular workflow threatens the system's stability, the YCRC may temporarily lock an account without prior notice, with account restoration requiring consultation with YCRC staff to address the workflow.

## Research Support and User Assistance

The YCRC includes several research support staff members who can help users with a variety of tasks, including education and training, software installation, and workflow assistance. The YCRC offers a number of classes and workshops on a variety of topics relevant to research computing. New users are encouraged to attend one of the introductory training workshops (or watch a recorded training) to learn about the HPC clusters and become familiar with the YCRC's standard operating procedures.

The YCRC procures, installs, and maintains many standard software tools and applications intended for YCRC facilities, including its HPC clusters. Among these are compilers and languages (e.g., Python, C, C++, Fortran), parallel computing tools (e.g., MPI), application systems (e.g., R, Matlab, Mathematica), and libraries (e.g., Intel Math Kernel Library, NAG, GNU Scientific Library, FFTW). Users requiring additional software on YCRC facilities are encouraged to install their own copies, though the YCRC's research support staff can assist as needed. For customizable systems such as Python and R, the YCRC has set up procedures to enable users to easily install their own modules, libraries, or packages.

The YCRC has set up its own "ticketing system" to help manage and address inquiries, requests, and troubleshooting issues related to the HPC clusters. Users may contact YCRC staff by emailing [research.computing@yale.edu](mailto:research.computing@yale.edu). While the time required to resolve particular issues may vary widely, users may expect an initial communication from a YCRC staff member within a reasonable time (often within one business day). Users may also obtain assistance by arranging appointments to meet with the research support staff or attending the [YCRC Office Hours](#).

## Hardware Acquisition and Lifecycle

At least once per year, the YCRC will offer PIs an opportunity to purchase dedicated HPC compute and storage resources for their groups. Often, such opportunities may be coordinated with the YCRC's regular refresh and upgrade cycles for its compute and storage hardware infrastructure, but, subject to YCRC approval, PIs may request hardware purchases at other times, as needed. At its discretion, the YCRC may restrict the types and quantities of dedicated HPC resources purchased for PI groups and require that purchased resources be compatible with the YCRC's data center and network infrastructure and with applicable policies.

### **Procurement Process & Approved Vendors**

All HPC resources will be procured through the YCRC from vendors of the YCRC's choosing and purchased with warranties acceptable to the YCRC (currently for 5 years). HPC resources will have lifetimes consistent with their warranties, commencing upon delivery. After which, the resources may be decommissioned, or the lifetimes may be extended, at the YCRC's sole discretion.

### **Hardware Decommission and Expiration**

All hardware purchased for the YCRC clusters includes a multi-year warranty and is subject to decommission any time after that date. For storage, and if the owner desires to retain the data, new storage must be purchased in advance of warranty expiration to ensure time to migrate retained data. Data hosted on storage that is not renewed will be purged and unrecoverable after the expiration date. Compute nodes will be supported by YCRC in conjunction with the vendor during the warranty and may be run beyond their warranty at the YCRC's discretion but may experience unrecoverable failure and should not be expected to remain in service for any amount of time after the warranty expires. The YCRC will decommission all nodes, whether or not they have failed, within two years of the expiration of their warranty to ensure capacity for new nodes and to maintain a modern HPC environment.

### **Private Partition and Resource Sharing**

Each PI group may access its dedicated compute resources using a private Slurm partition restricted to use by group members. The YCRC reserves the right to allow other users to use idle dedicated compute resources by submitting jobs to a Slurm scavenge partition. Any scavenge job will be subject to termination after a minimum run time (defined per cluster) should a job be submitted to the group's private partition requiring the node occupied by the scavenge job.