

Getting Started

If you read nothing else, read this.

This page is an insanely simple guide to just barely get you going. The documentation on this page will *not* be complete. It will *not* tell you everything. It *will* tell you just enough to get started with using Research Computing Cluster. First off, email rc-help@rit.edu if you encounter any issues with the cluster or if your account doesn't get created within 1 hour of completing the Apply process.

Apply for Access to Research Computing

Go to <http://rc.rit.edu/apply> and login with your RIT username and password. It's as simple as that! If 1 hour passes without receiving an account creation email, please contact us at rc-help@rit.edu.

Logging in to Research Computing

You may either use an SSH client or Remote Desktop (RDP) client to connect to `ion.rc.rit.edu`, the Research Computing Cluster Job Submission Node. When prompted for credentials, please use your RIT username and password.

Check the full articles on [SSH](#) and [RDP](#), if it is your first time working with remote systems. For more information on how to move files between your RC Home Directory and another machine, please see [Remote File Access](#).

SLURM – The Cluster Scheduler

Simple Linux Utility Resource Manager. It manages running your research computation on the Research Computing Cluster. You tell it what resources you need, it dispatches your work to one or more of those computers, and then the program runs to completion. When your work is done, SLURM frees up those resources so that someone else can use them. There are two ways to submit work:

Using a SLURM Batch File (sbatch)

This is the preferred method for creating work for the Research Computing Cluster. An SBATCH file consists of some commands telling SLURM what resources you need, followed by the Linux commands necessary to perform the requested work. Here you can tell SLURM how many cores (CPU) you need and how much memory (RAM), as well as how long to let your job run before it has taken too long. This file can then be run using the "sbatch" command to submit it to the SLURM Queue.

For more information on SLURM parameters see [SLURM Scheduler](#)

For example SBATCH files, please run "grab-examples" from the command line. This will create the folder "slurm-examples" in your Home Directory and fill it with some scripts showing you the basic usage of the cluster.

Using Interactive Mode (sinteractive)

If you need user interaction or are only running something once then run `sinteractive`. This will ask you for the resources you require and then connect you to the scheduled node. If you don't know what that entails, just try it. Be sure to exit from your interactive session when you're done, otherwise you're a terrible person for requesting resources you aren't using. See the full SLURM user documentation if you need to make batch jobs or want to learn about other cool commands like `squeue`, `sbatch`, `scontrol`, `sacct`, `scancel`, and `smap`.

Tips

- Try to avoid requesting more resources than you need on the cluster. Any resources that you tell SLURM you're going to use cannot be offered to other users regardless of whether you actually use those resources. If you fudge your numbers too much, you're eligible to have your allocation **terminated** by admins.
- "sinteractive" is not the best way to run jobs by any means. For any sort of repeatable workflow, you will need to write SBATCH jobs.
- "sinteractive" uses the Linux Screen utility. For more help see [Screen Help](#).

Software

Basic interpreted languages are there (Python, Perl, BASH). If you can't find something you need, we might have it as a module. Run ``module avail`` to see what modules we have available for loading. From that list, you can run ``module load <NAME>/<VERSION>`` where `<NAME>/<VERSION>` is from that list. For example, you can run ``module load R/3.1.1`` and then you magically have access to that version of R.

Need More Help?

If you need help using any Research Computing resources, don't hesitate to stop by our lab. Our infrastructure is changing all the time and our documentation may not always be up-to-date. We're more than happy to work with you to accommodate your research needs.