

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 the Research Computing cluster, SPORC (Scheduled Processing on Research Computing). First off, email rc-help@rit.edu if you encounter any issues with the cluster or if your account doesn't get created within one hour of completing the Apply process.

Research Computing also has [a new website!](#) It has much of the same information as the wiki, but also includes our questionnaire, a place to catalog your research projects with us, and is overall aesthetically prettier (if you love RIT branding, you'll love this site). We are still in the process of making it as useful for researchers as possible so expect changes, and if you have any feedback don't be hesitant to contact us.

If you are unfamiliar with the topic of research computing or want to know more about what we offer check out our [Services](#) and [About](#) page.

Questionnaire

Research Computing aims to know the researchers, their research, and their computational and research computing needs. After creating an account make sure to fill out [this questionnaire](#) so we can provide what you need for your project. After completion, your response can be found under [My Projects](#).

Logging in to Research Computing

Now that you have an account, you are ready to get onto our cluster, SPORC. To log into SPORC or other Research Computing provided resources, you will need to use either SSH or FastX. The first time you access SPORC, I would recommend you do it through SSH. FastX is useful for when you need a graphical interface, such as when you are visualizing data or analyzing graphs. When you log into SPORC there will be a message of the day that contains helpful tips and important information, so make sure you read it. You can view the message anytime by running `motd`. **Remember:** the host name for SPORC is `sporcsbmit.rc.rit.edu`.

[View SSH documentation.](#)

[View FastX documentation.](#)

Navigating the Terminal

Whether you logged in with SSH or FastX, it important that you know how to navigate the terminal. Through the terminal you will do most of the work for your research with Research Computing, including creating and editing files, loading software, and running jobs. If you have never used a terminal before, start by learning some [Basic Unix Commands](#) and how to use a [Text Editor](#).

Transferring Files

You must use a file transfer method to get files on and off your Research Computing account. Linux and Mac users have a built-in command that they can run, but every platform can take advantage of a software to do it for you. We currently provide documentation for using FileZilla and in the future we will be using Globus. However, these are not your only option; if you are logged onto a computer with a graphical interface via FastX you can use web applications like Dropbox.

[View File Transfer documentation.](#)

SLURM – The Cluster Scheduler

Slurm is the software we use to manage your research computation on SPORC. You tell it what resources you need ([read this if you need help determining what you need](#)), and it dispatches your work to one or more computers that make up SPORC, and then the program runs to completion. When your work is done, Slurm frees up those resources so that someone else can use them. Below we have outlined two ways to submit jobs. Before writing and submitting your first job, become familiar with how to view your jobs and the resources available on SPORC.

View [Slurm Basic Commands](#) documentation.

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 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 by running `exit` when you're done, otherwise you're a terrible person for requesting resources you aren't using. For the full process, see our documentation.

View [Submitting a job with sinteractive](#) documentation.

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.

Software

Your research more than likely needs software, whether it's for creating, analyzing, or visualizing data. So before you start researching, you need to learn how to use modules. Modules are the way Research Computing provides a lot of software without it all conflicting with each other. If you need it, we might have it. If we don't have it email us.

View [Using Modules](#) documentation.

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.