

Community Cluster

Quick Reference Card



Rosen Center for Advanced Computing

Quick Start

How do I log in?

Use `ssh -Y` to connect to `myclustername.rcac.purdue.edu`. When you log in, you connect to a shared front-end machine on the cluster.

How do I run my applications?

You must submit jobs to Slurm queues for running your applications, as described below. Applications run directly on the front-ends you log in to will be automatically terminated.

What is a job?

A job is essentially a request for compute resources. After logging in to a cluster, you typically submit a job to get an allocation on a compute node for a specific duration.

How do I submit a job?

You can submit a job by running appropriate Slurm commands. Jobs can either be: a) batch - you submit a script to Slurm, or b) interactive - you manually type the commands in a terminal.

What is Slurm?

Slurm is a software (also called the scheduler) that runs on the cluster and allocates resources to users in units of jobs.

What information do I need to provide when submitting a job?

You typically need to specify how many nodes and cores are needed by your application, for how long, and the queue to submit the job to.

What is a queue?

A queue is a resource pool for your research group which gives you priority (faster) access to the compute nodes.

Which queues can I use?

Use `slist` to see the names, status, and limits of all your queues.

Using Slurm

Q. How do I submit a batch job?

To submit the job submission file `myfile.sh` to the queue `myqueue` and request 1 node with 32 processor cores per node and 32 MPI ranks per node for a maximum run time of 5 hours, run the command:

```
sbatch -A myqueue -N 1 -n 32 -t 5:00:00 myfile.sh
```

Q. How do I submit an interactive job?

To submit an interactive job to the queue `myqueue` and request one core (for the default duration of 30 minutes), run the command:

```
sinteractive -A myqueue
```

Q. Can you provide an example job script?

Example job submission scripts can be found in the community cluster user guides.

<https://www.rcac.purdue.edu/knowledge>

Other Slurm Commands

`slist` List all queues that I can use and their current status and limits.

`sfeatures` List the hardware resources available on all node types.

`squeue -A myqueue` List all current jobs in the queue `myqueue`.

`squeue -u myusername` List all current jobs from the user `myusername`.

`jobinfo myjobid` View all information about the job ID `myjobid`.

`geany slurm-myjobid.out` View output from the currently running or completed job ID `myjobid`.

`scancel myjobid` Stop and delete the job ID `myjobid`.

`scontrol hold myjobid` Hold the job ID `myjobid` in the queue so it will not run.

`scontrol release myjobid` Release the held job ID `myjobid`, and allow it to run.

*Underlined words should be replaced with their actual values when running a command, e.g., `myqueue` should be replaced with your queue name.

Community Cluster

Quick Reference Card



Rosen Center for Advanced Computing

Storage

What storage options are available?

Name	Capacity	Purpose	Location	Environment variable	Lost file recovery	File transfer tools	Notes
Home	25 GB	Persistent data (codes, software, important results)	<code>/home/myusername</code>	\$HOME	Yes. Use "flost"	Globus, SCP, SFTP	<ul style="list-style-type: none"> Use "myquota" to check usage
Scratch	200 TB	Temporary data (datasets, temporary files, results from simulations)	<code>/scratch/clustername/myusername</code>	\$SRCAC_SCRATCH	No. Periodically archive your files in Fortress.	Globus, SCP, SFTP	<ul style="list-style-type: none"> Unused files purged every 60 days Use "purge" to see which files will be purged. Use "myquota" to check usage
Data Depot	On demand	Shared group storage for persistent data (software, datasets, results)	<code>/depot/mygroup</code>		Yes. Use "flost"	Globus, SCP, SFTP, Network Drive	<ul style="list-style-type: none"> Purchase required Use "myquota" to check usage
Fortress	Unlimited	Data archival (back up your files here)	<code>/home/myusername</code> <code>/group/mygroup</code>		No.	Globus, HSI, HTAR	<ul style="list-style-type: none"> It is recommended to combine small files into tar archives before transfer.

Graphical and Interactive Use

Q. How do I use a graphical application on the cluster?

There are multiple ways to launch a graphical application on the clusters as shown in the table below. The easiest solution will depend on which application you are trying to use and on your workflow.

	Thinlinc Desktop	Open OnDemand	VS Code	Interactive Job
How to connect	Navigate to <code>desktop.clustername.rcac.purdue.edu</code>	Navigate to <code>gateway.clustername.rcac.purdue.edu</code>	Connect to <code>loginXX.clustername.rcac.purdue.edu</code>	SSH to <code>clustername.rcac.purdue.edu</code> , then submit an interactive job
How to use	From the top menu go to "Applications" -> "Cluster Software"	From the top menu, go to "Interactive Applications" -> Fill out necessary details -> then submit	Follow your software development workflow.	Once your job has started, load the necessary modules, then launch the application. For example: <code>\$ module load matlab</code> <code>\$ matlab</code>
Popular applications	Matlab, Rstudio, Ansys, Abaqus, etc.	Jupyter Notebook, Rstudio, Desktop, etc.	Jupyter Notebook, Python scripts	Any graphical application can be launched.
Important notes	Use Thinlinc client software for better graphical experience.	Remember to select the right queue and walltime for your job.	<ul style="list-style-type: none"> Setup SSH keys for easier login. Connect to a specific front-end node. Connecting to compute nodes requires SSH tunnel via a front-end. 	Your local SSH client must have an X11 server for graphics support.

Need more support?

Email Support: rcac-help@purdue.edu

Please include which cluster you are using, your Slurm job ID, applications (or modules) you are using, and any error messages or output you have received.



Coffee hour consultation
www.rcac.purdue.edu/coffee



Online documentation
www.rcac.purdue.edu/knowledge