



Getting started on the cluster



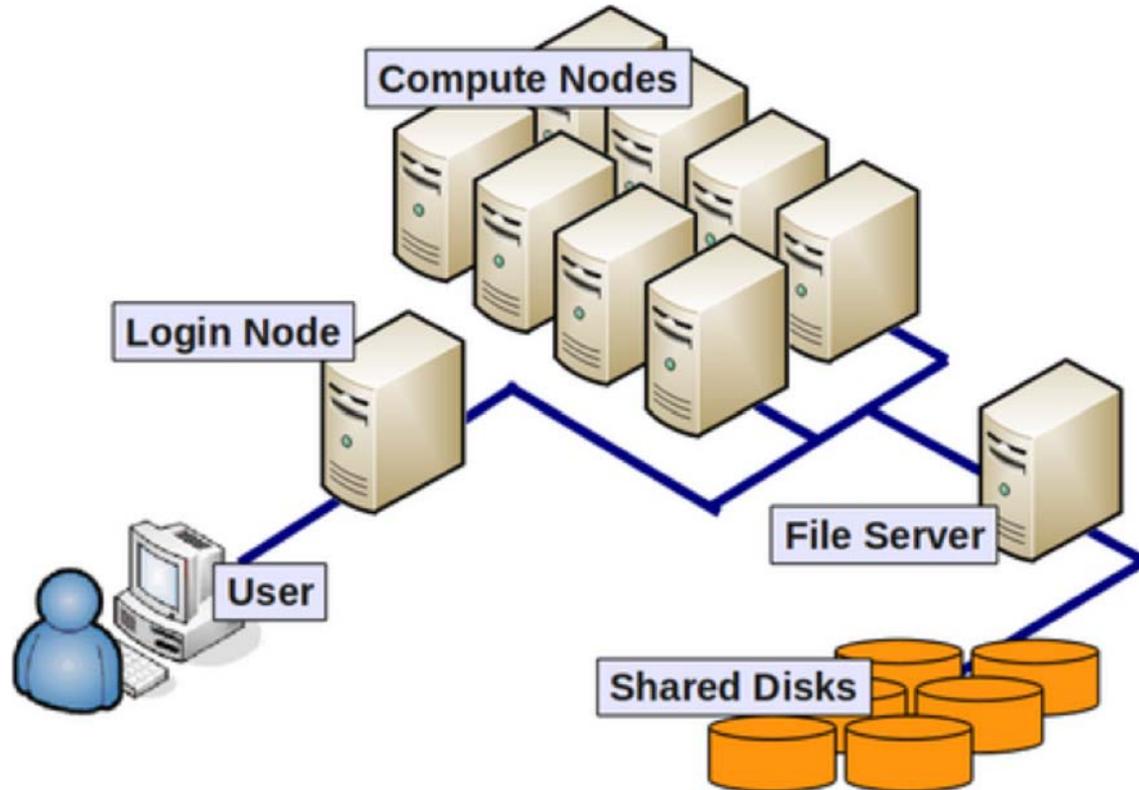
Learning Objectives

Describe the structure of a compute cluster

Log in to the cluster

Demonstrate how to start an interactive session with the SLURM job scheduler

Cluster Architecture





Cluster Terminology

- Supercomputer/High Performance Computing (HPC) cluster: A collection of similar computers connected by a high speed interconnect that can act in concert with each other
- Node: A computer in the cluster, an individual motherboard with CPU, memory, local hard drive
- CPU: Central Processing Unit, it can contain multiple computational cores (processors)
- Core: Basic unit of compute that runs a single instruction of code (a single process)
- GPGPU/GPU: General Purpose Graphics Processing Unit, a GPU designed for supercomputing.



Login & Access

<https://docs.rc.fas.harvard.edu/kb/quickstart-guide/>

Cluster Quick Start Guide

[Table of Contents](#) > [\[show\]](#)

This guide will provide you with the basic information needed to get up and running on the FASRC cluster for simple command line access. If you'd like more detailed information, each section has a link to fuller documentation

PREREQUISITES

1. Get a FASRC account using the account request tool.

Before you can access the cluster you need to request a Research Computing account.

See [How Do I Get a Research Computing Account](#) for instructions if you do not yet have an account.

See the account confirmation email for instructions on [setting your password](#) and getting started.

Login & Access

Once you have an account you can use the Terminal to connect to the cluster



– Mac: Terminal



– Linux: Xterm or Terminal



– Windows: SSH client - Putty or Bash Emulator - Git Bash

```
$ ssh username@login.rc.fas.harvard.edu
```

- ssh stands for Secure SHell
- ssh is a protocol for data transfer that is secure, i.e the data is encrypted as it travels between your computer and the cluster (remote computer)
- Commonly used commands that use the ssh protocol for data transfer are, scp and sftp

Login & Access

Once you have an account you can use the Terminal to connect to the cluster



– Mac: Terminal



– Linux: Xterm or Terminal



– Windows: SSH client - Putty or Bash Emulator - Git Bash

```
$ ssh username@login.rc.fas.harvard.edu
```

Login issues? See <https://rc.fas.harvard.edu/resources/support/>

Password:

Verification code:

Login & Access

<https://www.rc.fas.harvard.edu/resources/quickstart-guide/>

Once you have run the ssh command:

- Enter your password (*cursor won't move!*)
- Add a verification code (2-Factor Authentication)

2. Setup OpenAuth for two factor authentication

Once you have your new FASRC account, you will need to set up our OpenAuth tool for two-factor authentication.

See the [OpenAuth Guide](#) for instructions if you have not yet set up OpenAuth.

For troubleshooting issues you might have, please see our [troubleshooting page](#).



OpenAuth is 2-factor authentication separate from HarvardKey and updates the token every 30 seconds

Login & Access

```
rsk394 -- rkhetani@holylogin03:~ -- ssh rkhetani@login.rc.fas.harvard.edu -- 92x40

!!!!!!!!!!!!!!!!!!!!!!!!!!!! Cannon !!!!!!!!!!!!!!!!!!!!!!!!!!!!!

Welcome to Cannon, a HPC resource for the research community,
hosted by Research Computing at HU's Faculty of Arts and Sciences.

+----- Helpful Documentation: -----+
| https://rc.fas.harvard.edu/resources/quickstart-guide/ |
| https://rc.fas.harvard.edu/running-jobs/              |
| https://rc.fas.harvard.edu/convenient-slurm-commands/ |
+-----+

+----- NEWS & UPDATES: -----+
+ OFFICE HOURS: Wednesdays noon-3pm, 38 Oxford, ROOM 100 (1st Floor conf room) +
+ Check our consulting calendar at: https://www.rc.fas.harvard.edu/consulting-calendar/ +
+ Check our training schedule at: https://www.rc.fas.harvard.edu/upcoming-training/ +
+-----+

NEXT MAINTENANCE: NOVEMBER 4TH 7-11AM

https://www.rc.fas.harvard.edu/maintenance

CANNON: Cannon is live! See the Running Jobs page for information about
the updated partitions.

https://www.rc.fas.harvard.edu/resources/running-jobs/#Slurm_partitions

For more about the new cluster see:

https://www.rc.fas.harvard.edu/fasrc-cluster-refresh-2019/

GENERAL: The general partition has been decommissioned. Please use
the shared partition. For high memory jobs use bigmem.

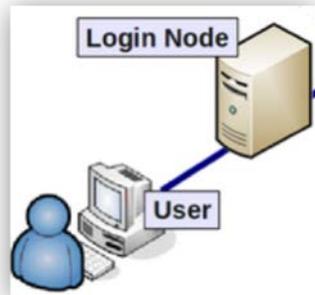
WINTER MAINTENANCE DECEMBER 3RD 7AM-5PM: We are doing an all day major
maintenance on December 3rd which will involve all running jobs being
cancelled. More details forthcoming soon. Please plan accordingly.

[ rkhetani@holylogin03 ~ ] $
```

You have logged into the login node!

```
[joesmith@holylogin03 ~]$
```

Name of the login node assigned to you





Access to resources on a compute node

- Login node:
 - not designed for analysis
 - not anything compute- or memory-intensive
 - best practice is to request a compute node as soon as you log in
- Interactive session:
 - work on a compute node “interactively”
 - request resources from SLURM using the `srun --pty` command
 - session will only last as long as the remote connection is active



Access to resources on a compute node

Simple Linux Utility for Resource Management - SLURM job scheduler:

- Fairly allocates access to resources to users on compute nodes
- Manages a queue of pending jobs; ensures that no single user or group monopolizes the cluster
- Ensures users do not exceed their resource requests
- Provides a framework for starting, executing, and monitoring batch jobs

Access to resources on a compute node

Requesting an interactive session:

```
[joesmith@holylogin03 ~]$ srun --pty -p test --mem 100 -t 0-01:00 /bin/bash
```

`srun --pty` - is how interactive sessions are started with SLURM

`-p test` - requesting a compute node in a specific partition*

`--mem 100` - memory requested in MB

`-t 0-1:00` - time requested (1 hour)

** Partitions are groups of computers that are designated to perform specific types of computing. More on next slide*

```
[joesmith@holy7c26602 ~]$
```

Name of the compute node
assigned to you



Partitions on the cluster

Partitions:	shared	gpu	test	gpu_test	serial_requeue	gpu_requeue	bigmem	unrestricted	pi_lab
Time Limit	7 days	7 days	8 hrs	1 hrs	7 days	7 days	no limit	no limit	varies
# Nodes	530	15	16	1	1930	155	6	8	varies
# Cores / Node	48	32 + 4 V100	48	32 + 4 V100	varies	varies	64	64	varies
Memory / Node (GB)	196	375	196	375	varies	varies	512	256	varies

Learn more about a partition:

```
$ sinfo -p shared
$ scontrol show partition shared
```

Request Help - Resources

<https://docs.rc.fas.harvard.edu/support/>

- Documentation
 - <https://docs.rc.fas.harvard.edu/documentation/>
- Portal
 - http://portal.rc.fas.harvard.edu/rcrt/submit_ticket
- Email
 - rchelp@rc.fas.harvard.edu
- Office Hours
 - Wednesday noon-3pm 38 Oxford - 100
- Consulting Calendar
 - <https://www.rc.fas.harvard.edu/consulting-calendar/>
- Training
 - <https://www.rc.fas.harvard.edu/upcoming-training/>

