UAB The University of Alabama at Birmingham

Getting Started with Cheaha: Basics for Your Research

Using UAB Research Computing’s Cheaha HPC system

Fortune Iriaye

Scientist I - Research Facilitator

Research Computing

December 2025

Outline

1. Introduction
2. Cheaha Access
3. Running and Submitting Jobs
4. Research Computing Services
5. Good Practices
6. Getting help and resources

What is Cheaha?

- Named after Mount Cheaha (highest natural point in Alabama)
- Cheaha is UAB’s (HPC) cluster managed by UAB IT Research Computing.
- Currently supports ~14000 CPU cores, ~100 NVIDIA-A100, and ~ 72NVIDIA-P100 GPUs.
- Designed to accelerate the completion of research compute workloads in significantly shorter times.

Who can use Cheaha?

- Individuals affiliated with UAB (Researchers, Faculty, Students)
- External Collaborators using XIAS Accounts. (require a UAB sponsor)

How do I access Cheaha?

Visit https://rc.uab.edu

Log in through SSO using your **BlazerID** and password

New? Then you'll fill out a short form to create your account.

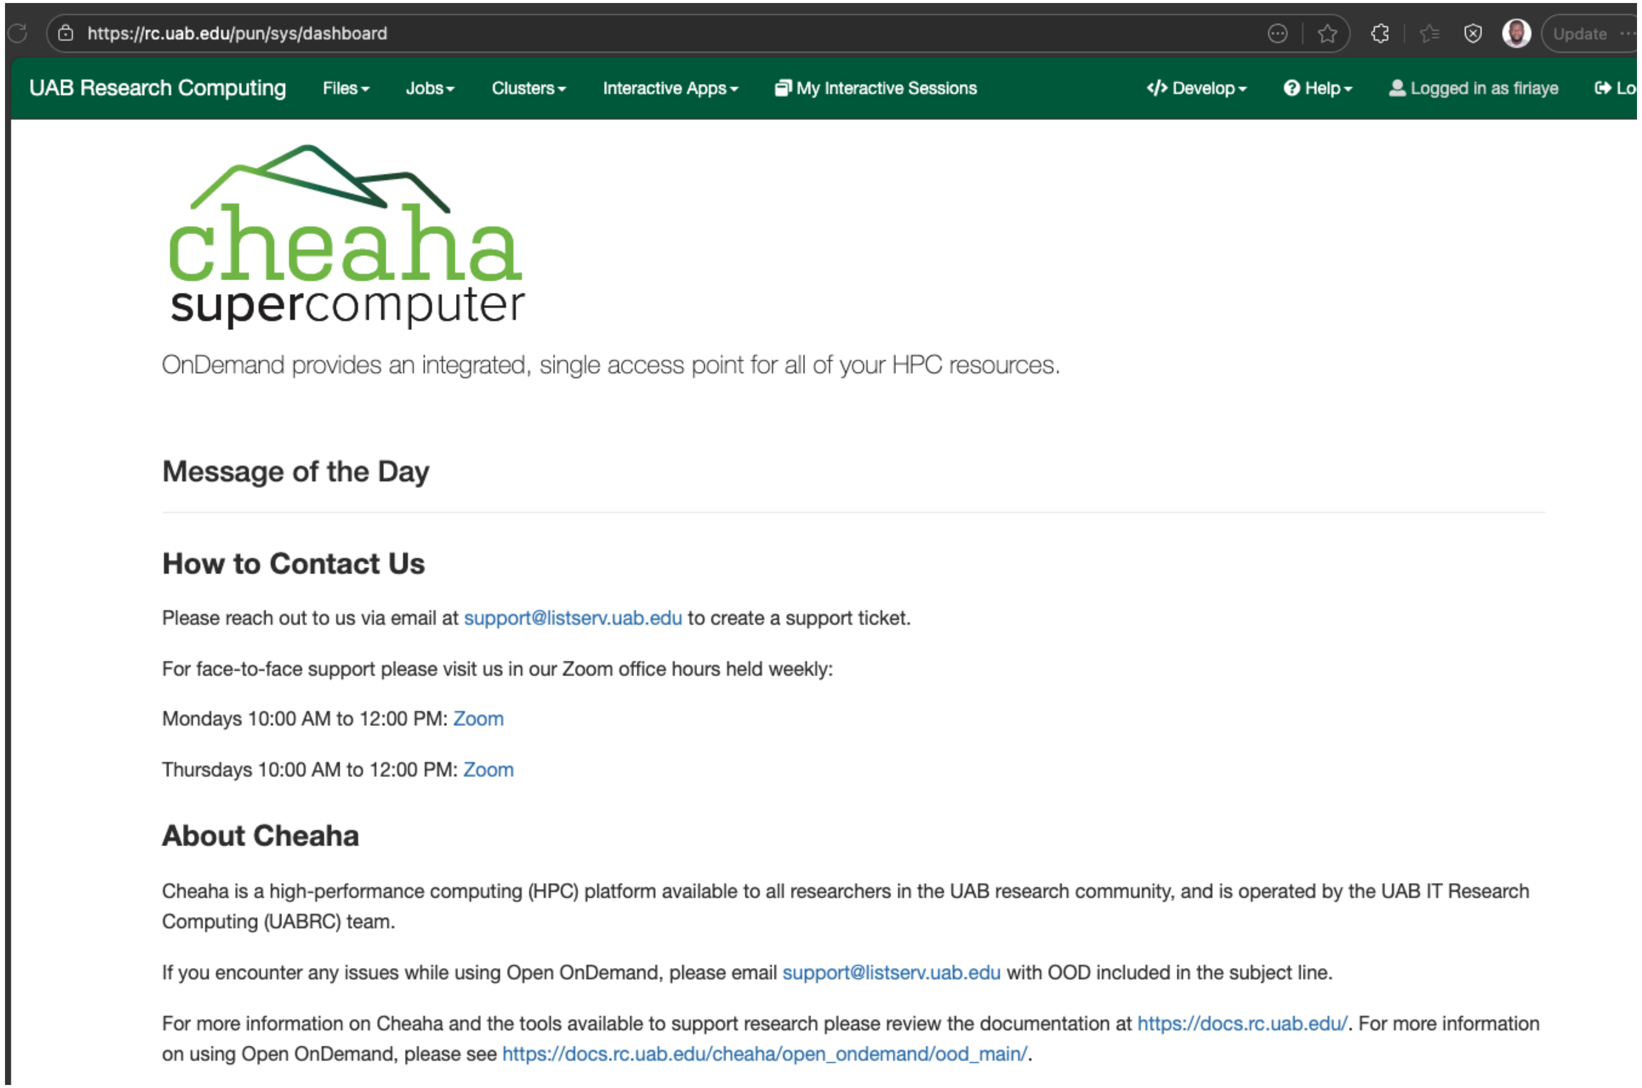

cheaha supercomputer

OnDemand provides an integrated, single access point for all of your HPC resources.

Message of the Day

How to Contact Us

Please reach out to us via email at support@listserv.uab.edu to create a support ticket.

For face-to-face support please visit us in our Zoom office hours held weekly:

Mondays 10:00 AM to 12:00 PM: Zoom

Thursdays 10:00 AM to 12:00 PM: Zoom

About Cheaha

Cheaha is a high-performance computing (HPC) platform available to all researchers in the UAB research community, and is operated by the UAB IT Research Computing (UABRC) team.

If you encounter any issues while using Open OnDemand, please email support@listserv.uab.edu with OOD included in the subject line.

For more information on Cheaha and the tools available to support research please review the documentation at https://docs.rc.uab.edu/. For more information on using Open OnDemand, please see https://docs.rc.uab.edu/cheaha/open_ondemand/ood_main/.

OPEN OnDemand

Running Jobs on Cheaha

Interactive Jobs via OOD, Batch Jobs with Slurm

Using Interactive Apps

Real-time, live work and code development

Great for testing and debugging

Ideal for Exploratory Data Analysis

Comfortable with Graphical User Interfaces (GUI)

How do I access and use Interactive Apps on Cheaha?

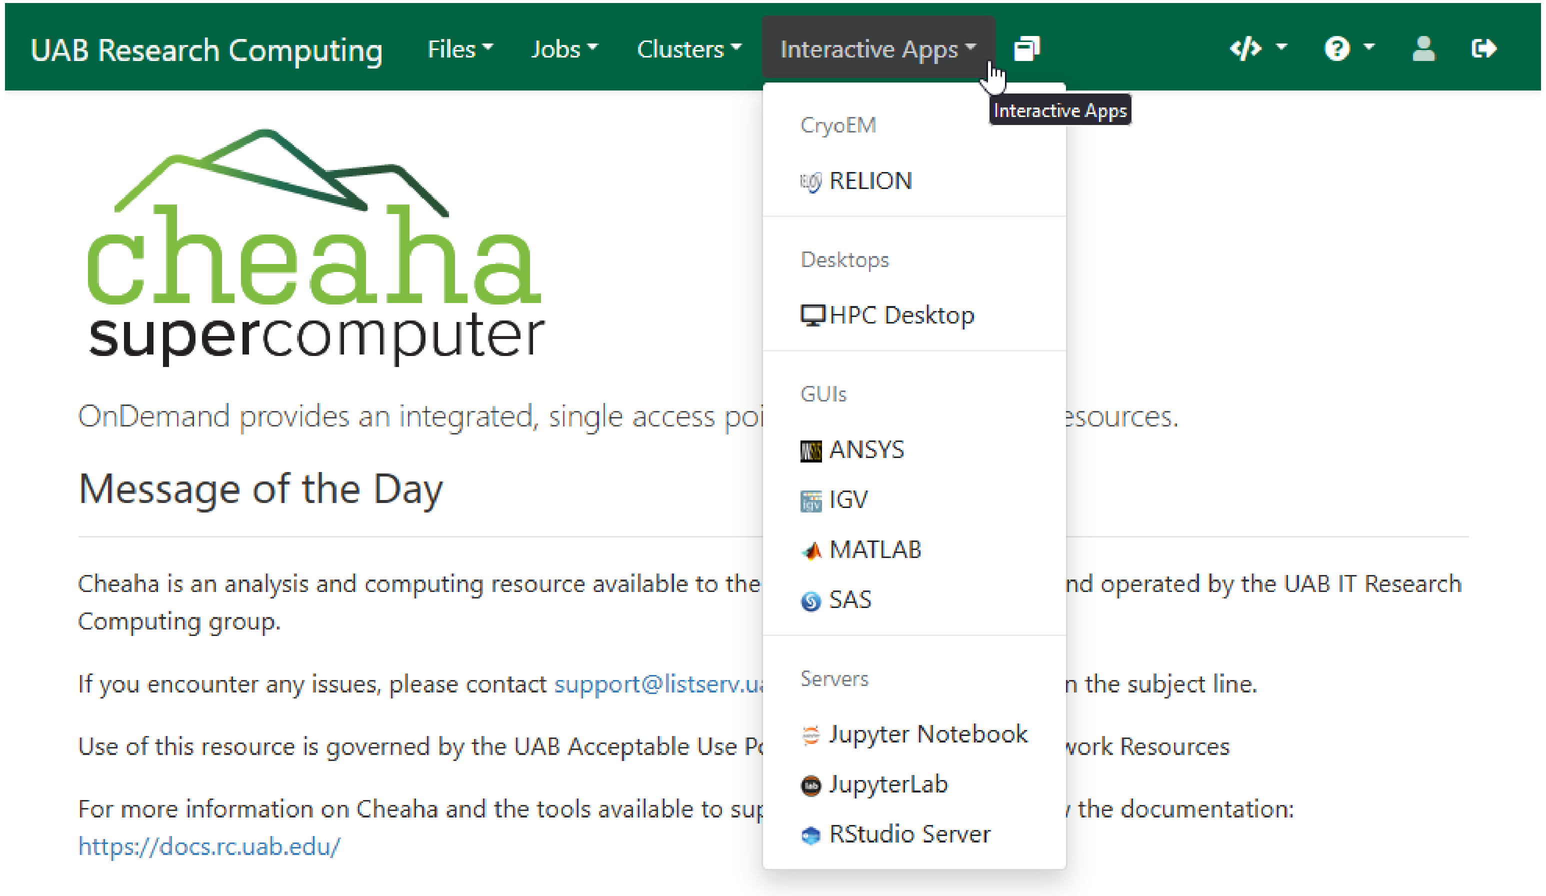

OnDemand provides an integrated, single access poi esources.

Message of the Day

Cheaha is an analysis and computing resource available to the nd operated by the UAB IT Research Computing group.

If you encounter any issues, please contact support@listserv.ua n the subject line.

Use of this resource is governed by the UAB Acceptable Use Po work Resources

For more information on Cheaha and the tools available to sup the documentation: https://docs.rc.uab.edu/

Click “Interactive Apps”…

How do I access and use Interactive Apps on Cheaha?

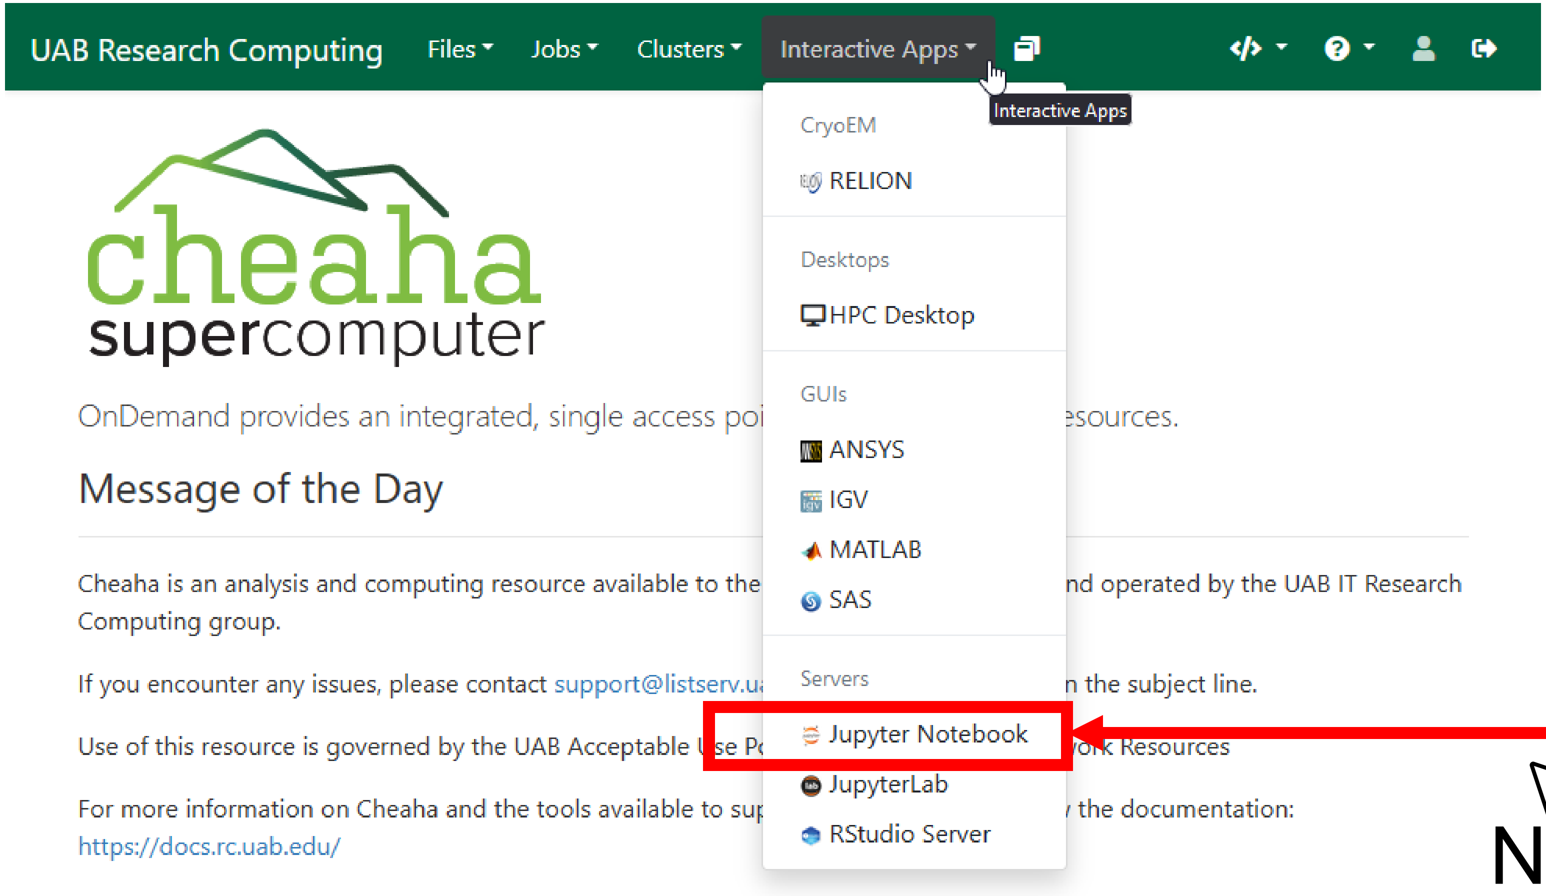

OnDemand provides an integrated, single access poi esources.

Message of the Day

Cheaha is an analysis and computing resource available to the nd operated by the UAB IT Research Computing group.

If you encounter any issues, please contact support@listserv.ua n the subject line.

Use of this resource is governed by the UAB Acceptable Use Po ork Resources

For more information on Cheaha and the tools available to sup the documentation: https://docs.rc.uab.edu/

Click “Interactive Apps”...

Click “Jupyter Notebook”

How do I access and use Interactive Apps on Cheaha?

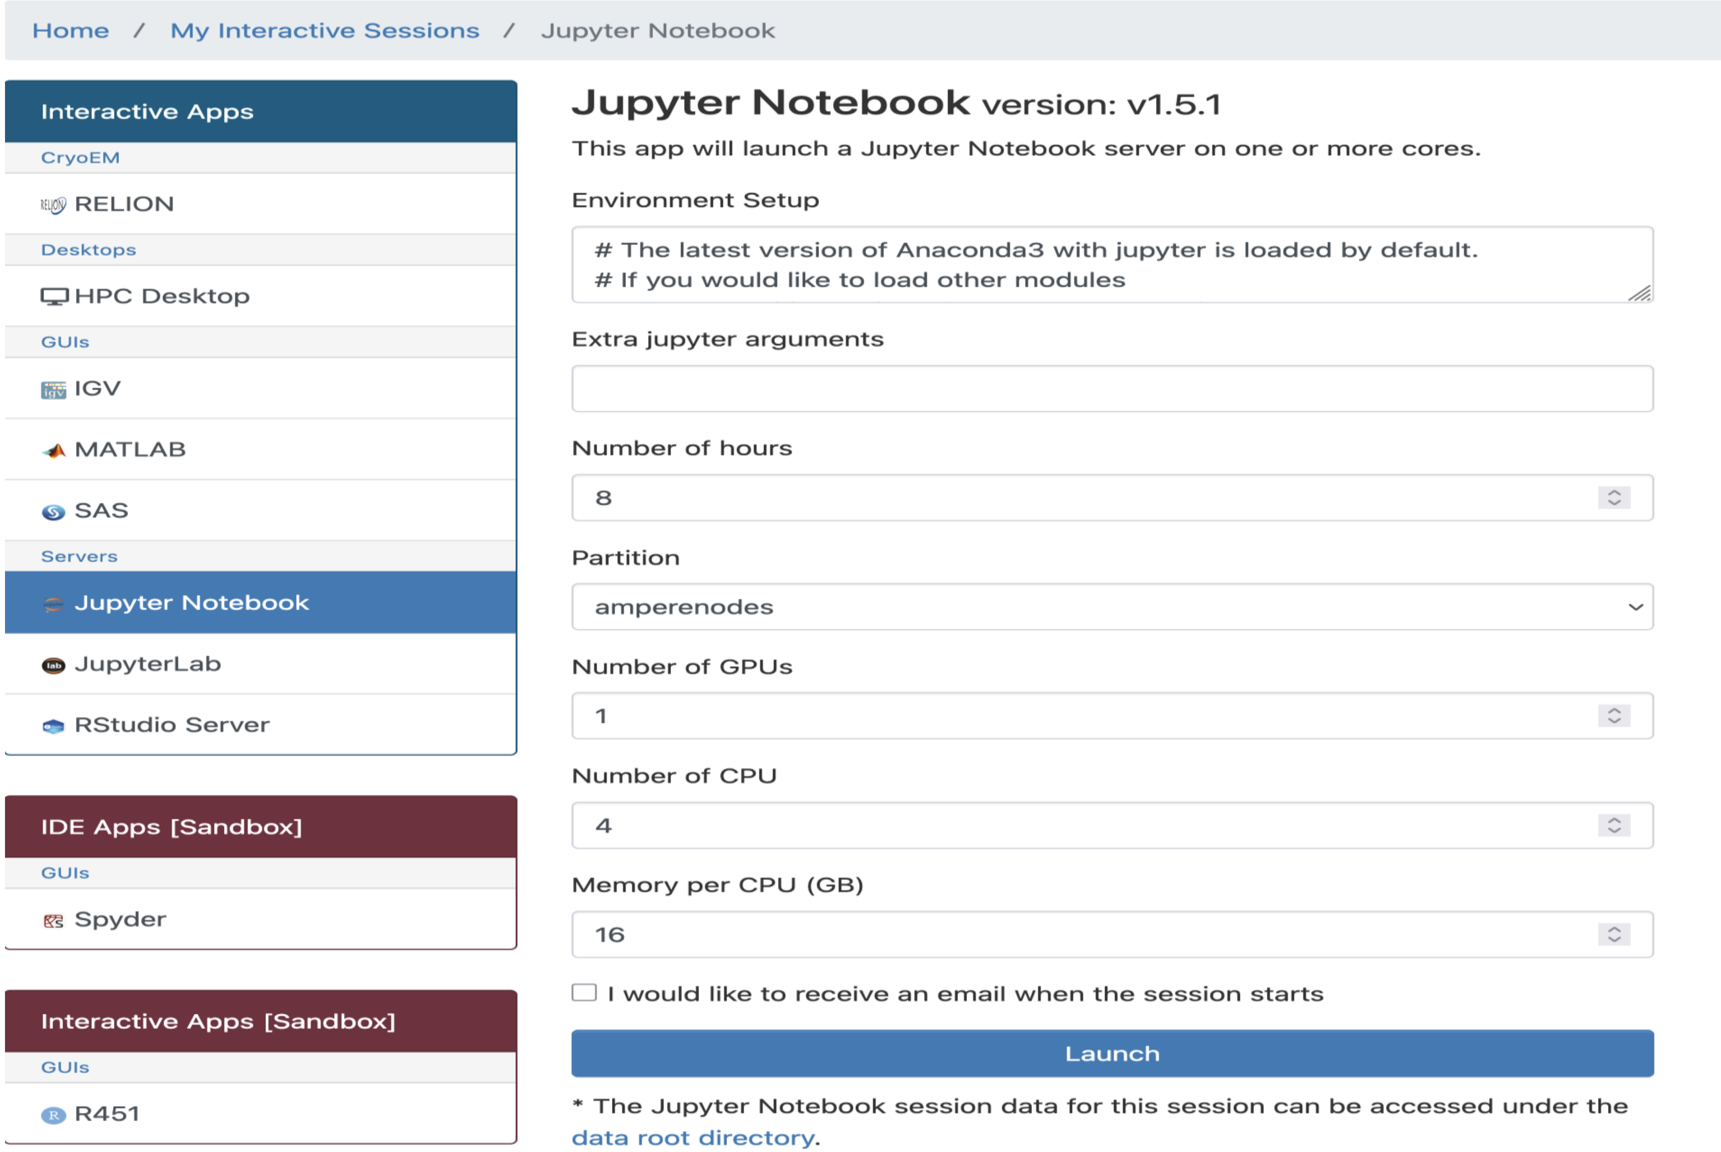

Jupyter Notebook version: v1.5.1

This app will launch a Jupyter Notebook server on one or more cores.

Environment Setup

```
# The latest version of Anaconda3 with jupyter is loaded by default.
# If you would like to load other modules
```

Extra jupyter arguments

Number of hours: 8

Partition: amperenodes

Number of GPUs: 1

Number of CPU: 4

Memory per CPU (GB): 16

☐ I would like to receive an email when the session starts

* The Jupyter Notebook session data for this session can be accessed under the data root directory.

The image shows a sample form

Select resources in the form

Click "Submit"

How do I access and use Interactive Apps on Cheaha?

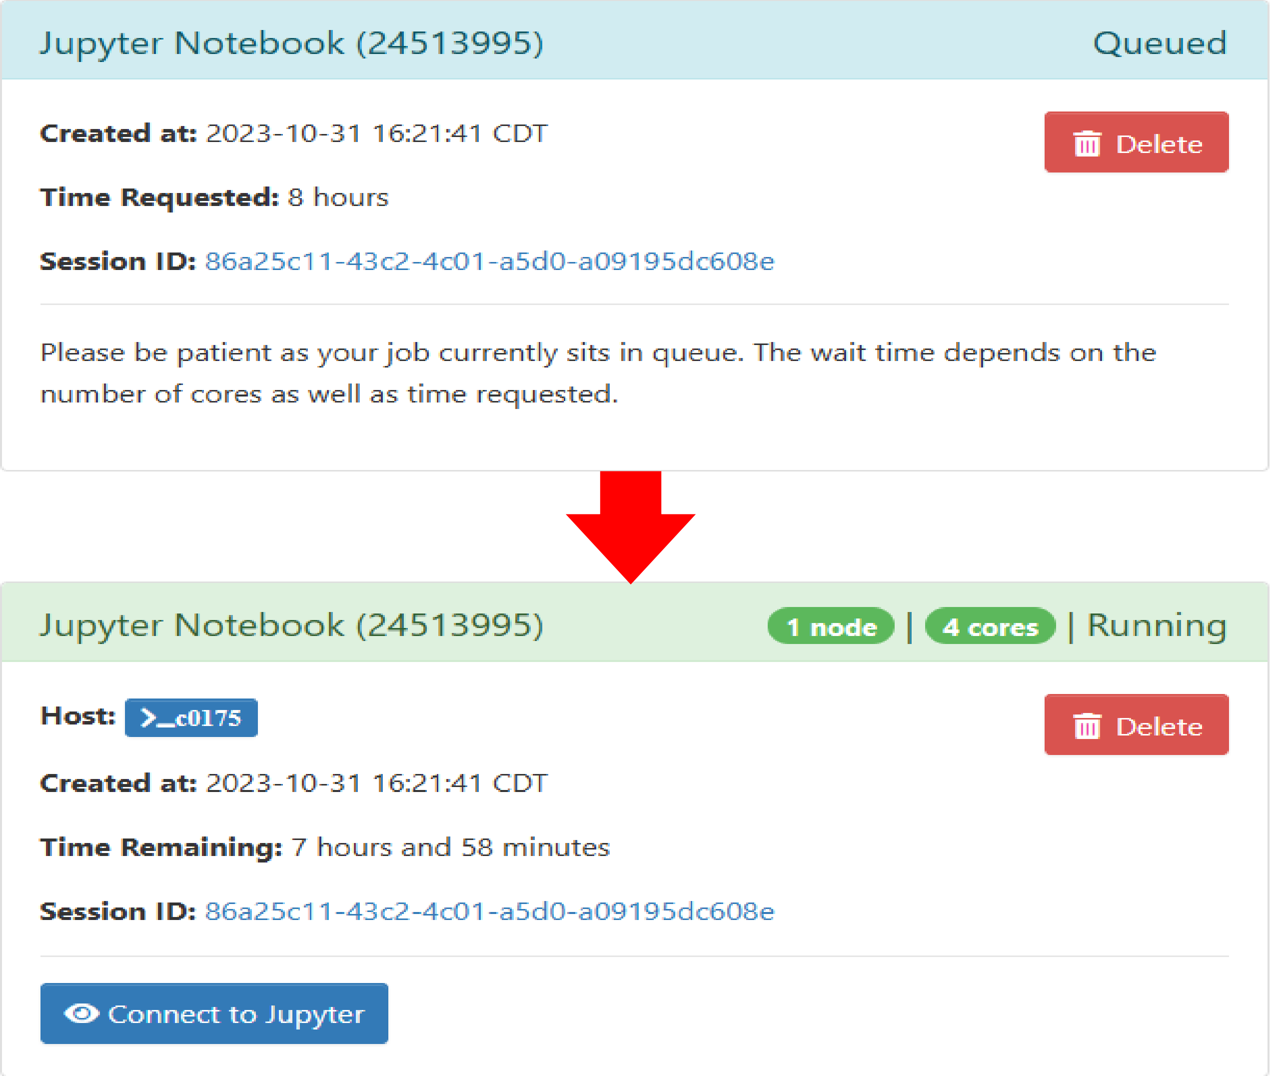

Jupyter Notebook (24513995) Queued

Created at: 2023-10-31 16:21:41 CDT

Time Requested: 8 hours

Session ID: 86a25c11-43c2-4c01-a5d0-a09195dc608e

Please be patient as your job currently sits in queue. The wait time depends on the number of cores as well as time requested.

Jupyter Notebook (24513995) 1 node | 4 cores | Running

Host: >_c0175

Created at: 2023-10-31 16:21:41 CDT

Time Remaining: 7 hours and 58 minutes

Session ID: 86a25c11-43c2-4c01-a5d0-a09195dc608e

- Wait for the job to start…

- Then click “Connect to Jupyter”

How do I access and use Interactive Apps on Cheaha?

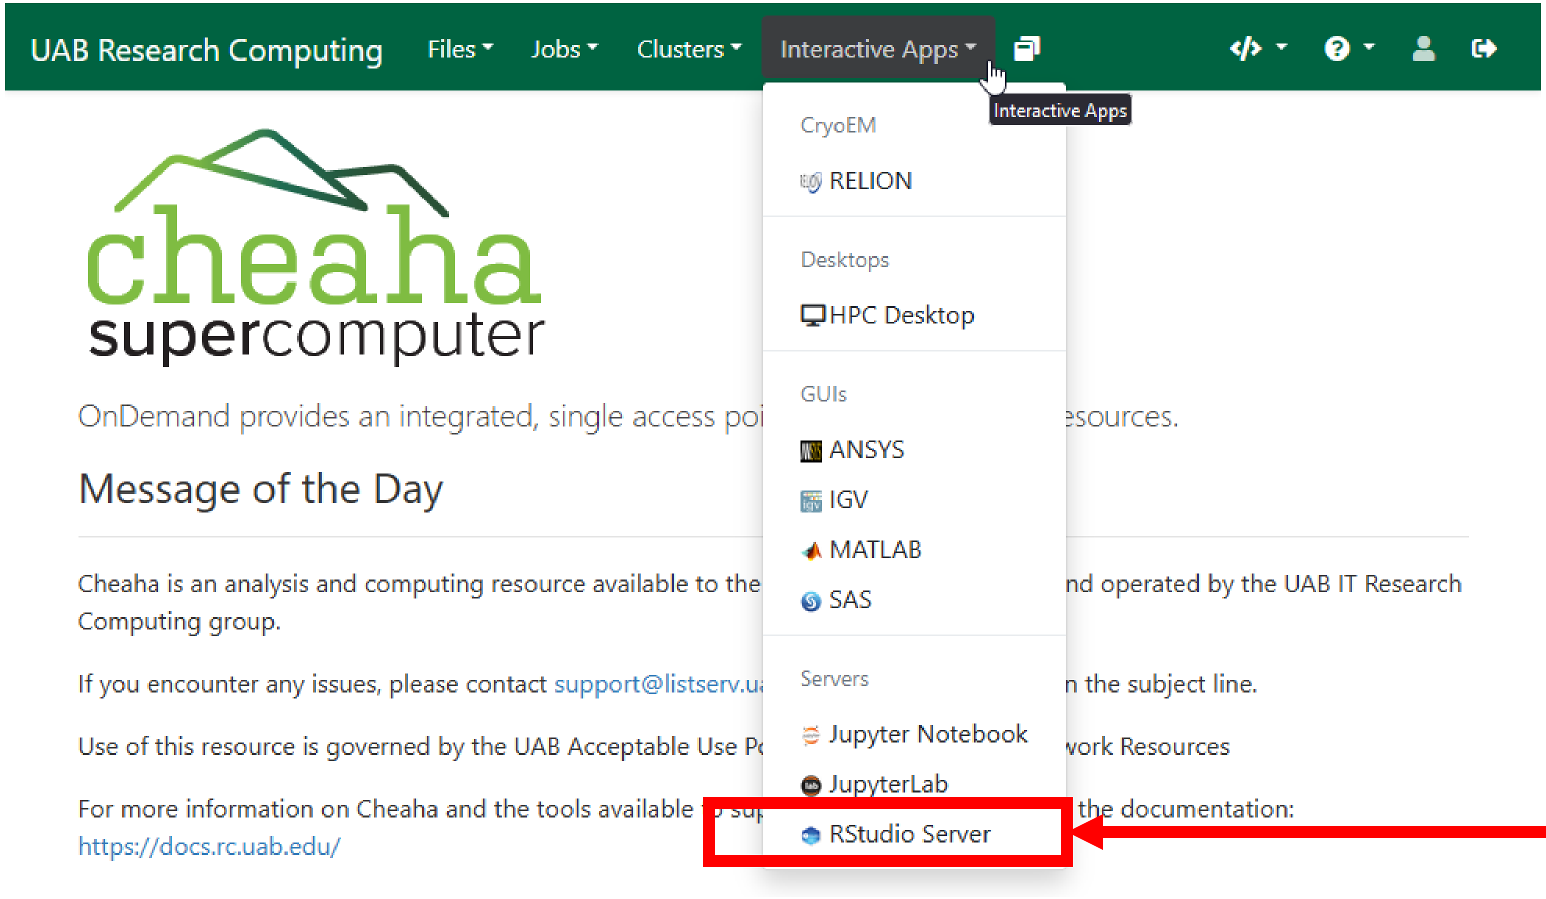

OnDemand provides an integrated, single access poi esources.

Message of the Day

Cheaha is an analysis and computing resource available to the nd operated by the UAB IT Research Computing group.

If you encounter any issues, please contact support@listserv.u n the subject line.

Use of this resource is governed by the UAB Acceptable Use Po ork Resources

For more information on Cheaha and the tools available o sup the documentation: https://docs.rc.uab.edu/

Click “Interactive Apps”…

Click “RStudio Server”…

OPEN OnDemand

How do I access and use Interactive Apps on Cheaha?

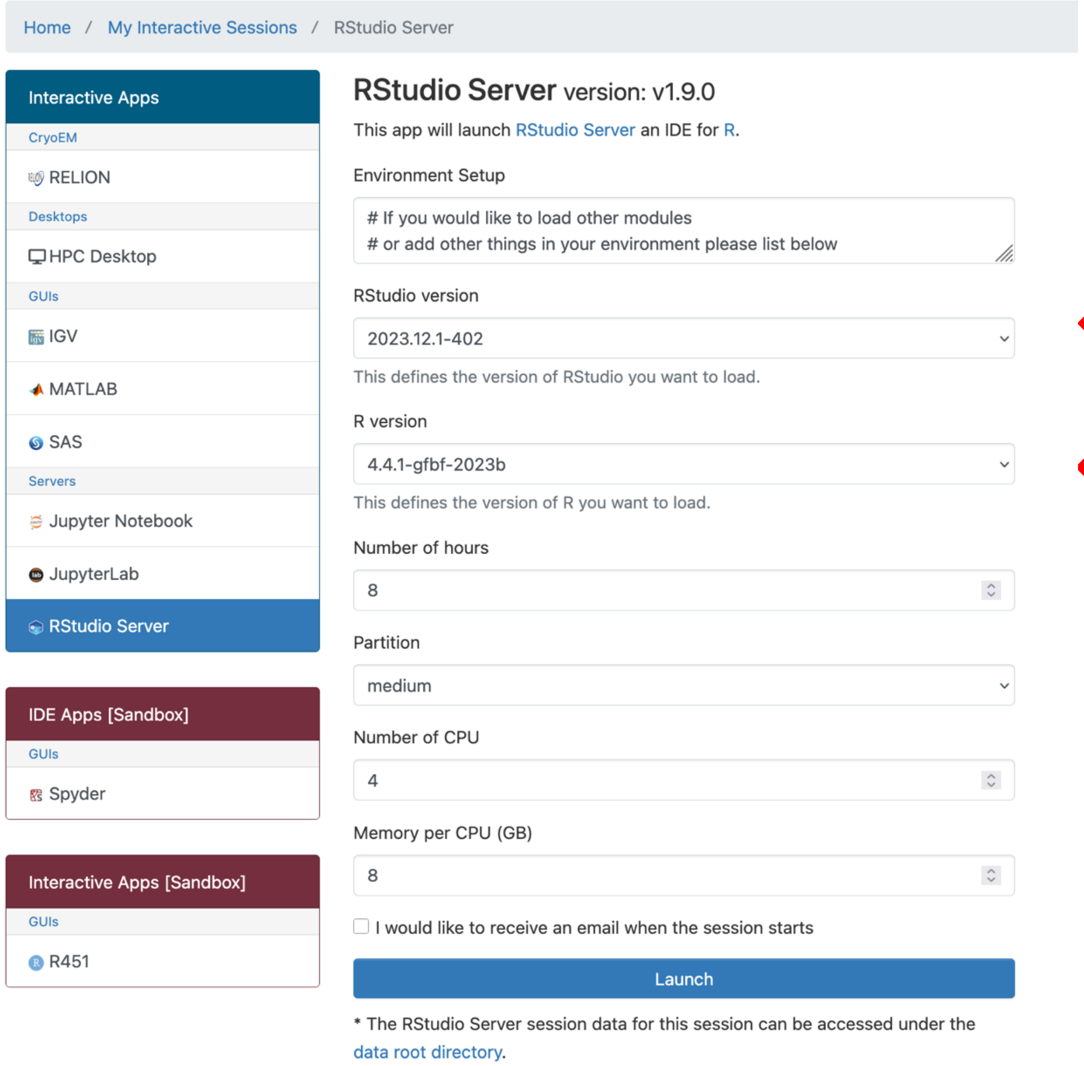

RStudio Server version: v1.9.0

This app will launch RStudio Server an IDE for R.

Environment Setup

```
# If you would like to load other modules
# or add other things in your environment please list below
```

RStudio version: 2023.12.1-402

This defines the version of RStudio you want to load.

R version: 4.4.1-gfbf-2023b

This defines the version of R you want to load.

Number of hours: 8

Partition: medium

Number of CPU: 4

Memory per CPU (GB): 8

☐ I would like to receive an email when the session starts

* The RStudio Server session data for this session can be accessed under the data root directory.

- Repeat the same process as before
- Fill the form relevant to your computational needs
- Note the addition of two fields for selecting RStudio version and R language version.

How do I access and use Interactive Apps on Cheaha (Desktop and Shell) ?

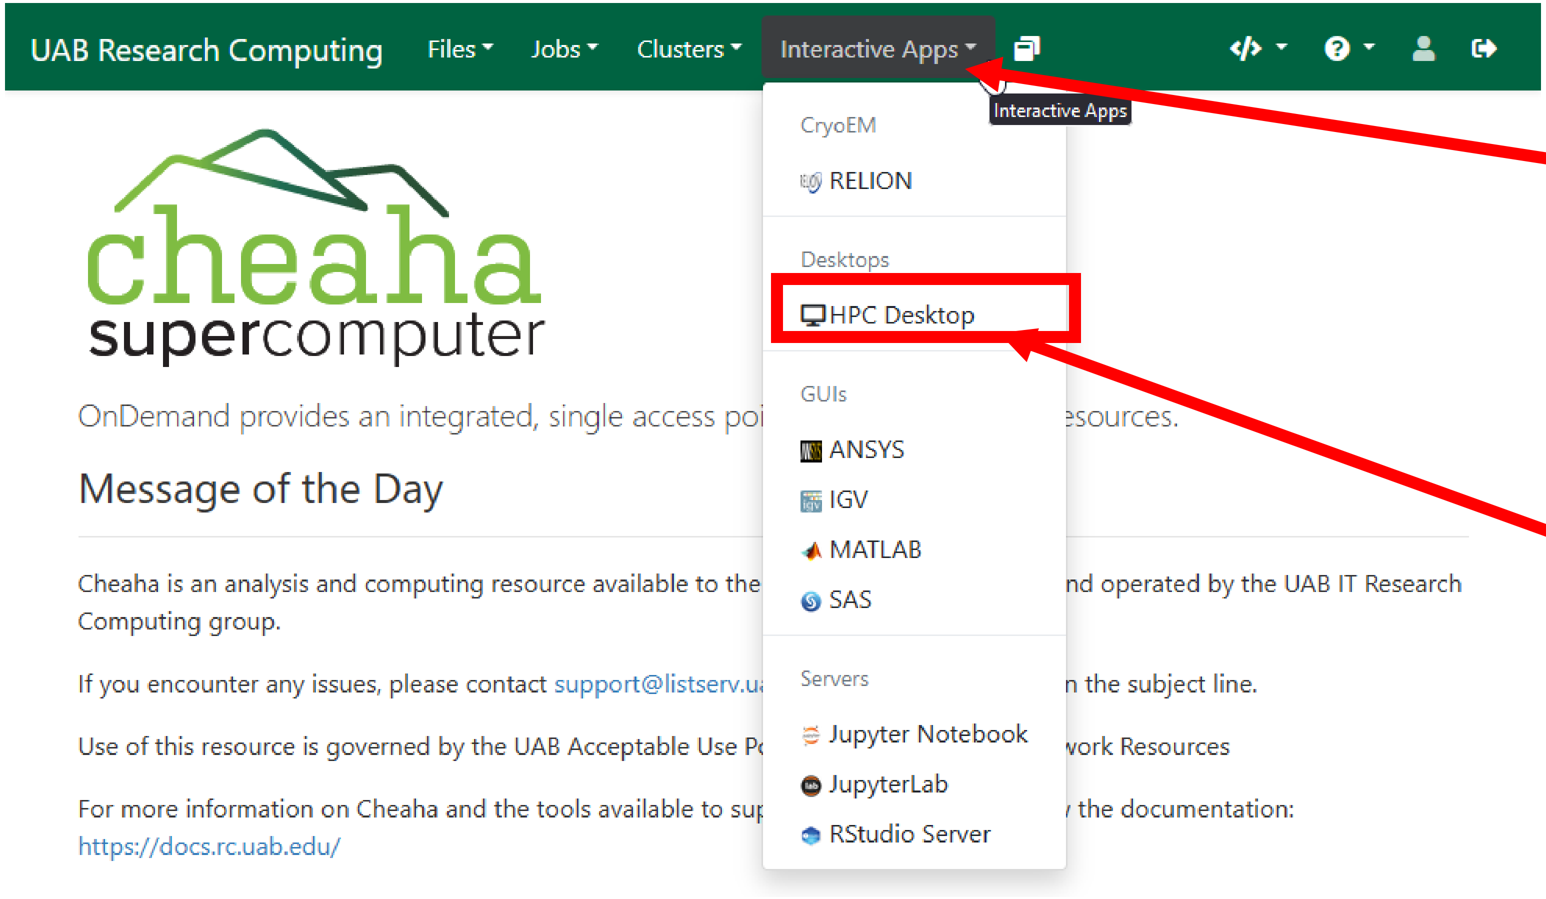

OnDemand provides an integrated, single access poi esources.

Message of the Day

Cheaha is an analysis and computing resource available to the nd operated by the UAB IT Research Computing group.

If you encounter any issues, please contact support@listserv.u n the subject line.

Use of this resource is governed by the UAB Acceptable Use P work Resources

For more information on Cheaha and the tools available to su the documentation: https://docs.rc.uab.edu/

Click “Interactive Apps”…

Click “HPC Desktop”…

OPEN OnDemand

Accessing Modules and Software Installed on Cheaha

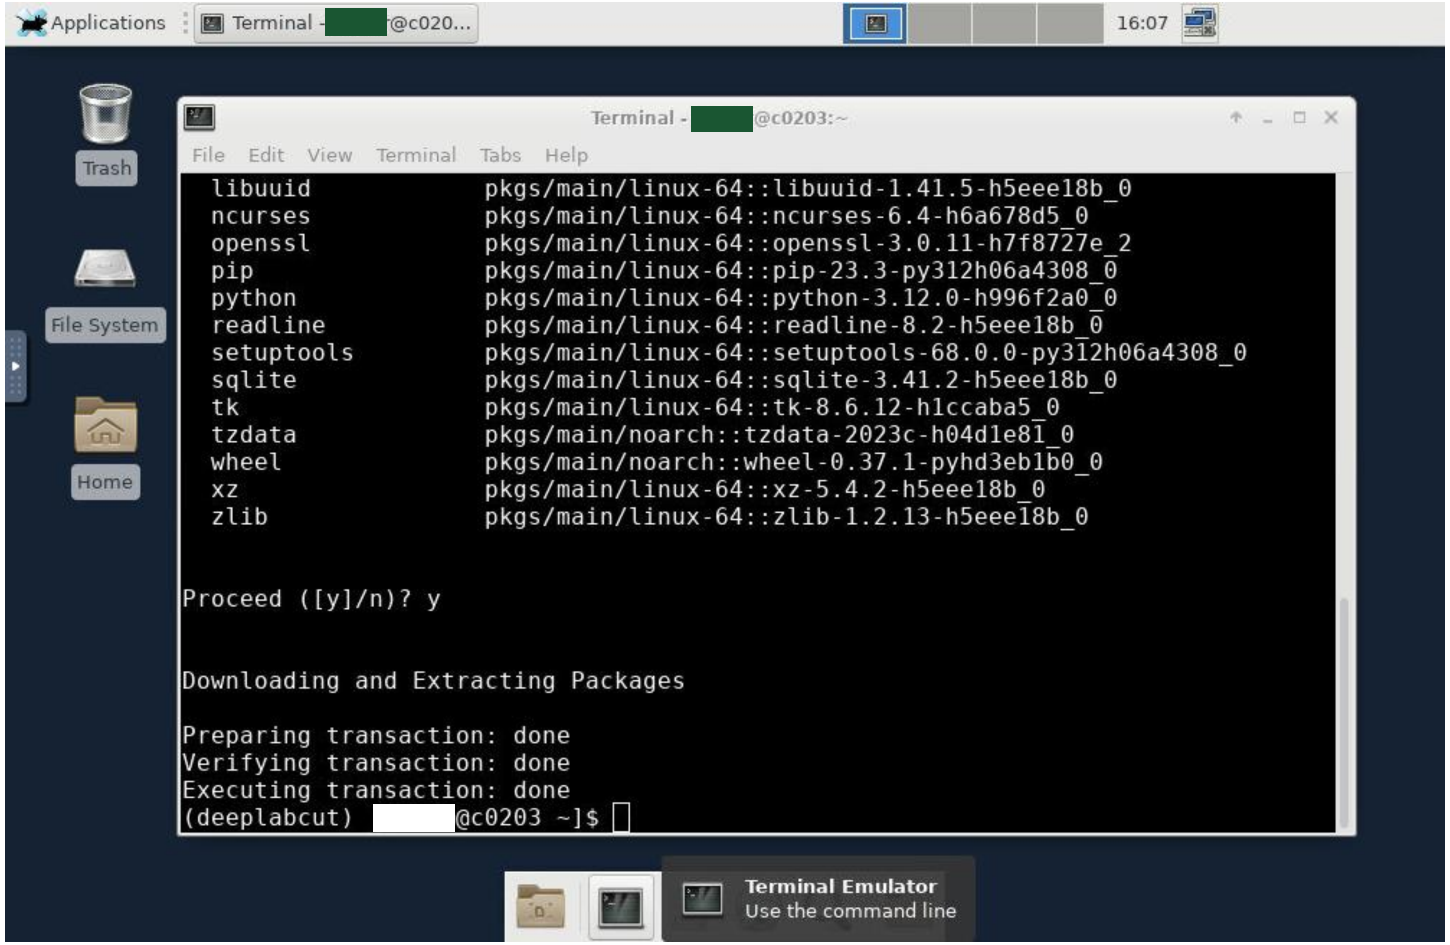

```
  libuuid            pkgs/main/linux-64::libuuid-1.41.5-h5eee18b_0
  ncurses            pkgs/main/linux-64::ncurses-6.4-h6a678d5_0
  openssl            pkgs/main/linux-64::openssl-3.0.11-h7f8727e_2
  pip                pkgs/main/linux-64::pip-23.3-py312h06a4308_0
  python             pkgs/main/linux-64::python-3.12.0-h996f2a0_0
  readline           pkgs/main/linux-64::readline-8.2-h5eee18b_0
  setuptools         pkgs/main/linux-64::setuptools-68.0.0-py312h06a4308_0
  sqlite             pkgs/main/linux-64::sqlite-3.41.2-h5eee18b_0
  tk                 pkgs/main/linux-64::tk-8.6.12-h1ccaba5_0
  tzdata             pkgs/main/noarch::tzdata-2023c-h04d1e81_0
  wheel              pkgs/main/noarch::wheel-0.37.1-pyhd3eb1b0_0
  xz                 pkgs/main/linux-64::xz-5.4.2-h5eee18b_0
  zlib               pkgs/main/linux-64::zlib-1.2.13-h5eee18b_0


Proceed ([y]/n)? y


Downloading and Extracting Packages

Preparing transaction: done
Verifying transaction: done
Executing transaction: done
(deeplabcut) [redacted]@c0203 ~]$
```

- Open the terminal with the icon shown at the bottom of the display.
- Load modules, environments, run scripts, etc.

Finding the Tool You Need

- Search a module you need using: *module spider <module name>*
- If you don't find It
  - The module may **not be installed system wide.**
  - Install it yourself, for example:
    - Python: consider using a Conda environment.
    - R: Install packages to a personal library.
    - Follow custom instructions for the tool.
- If a package is not available on CRAN or Conda:
  - Use containerized options: Pull images from Docker Hub and run with **singularity** on Cheaha.

Conda packages for Python

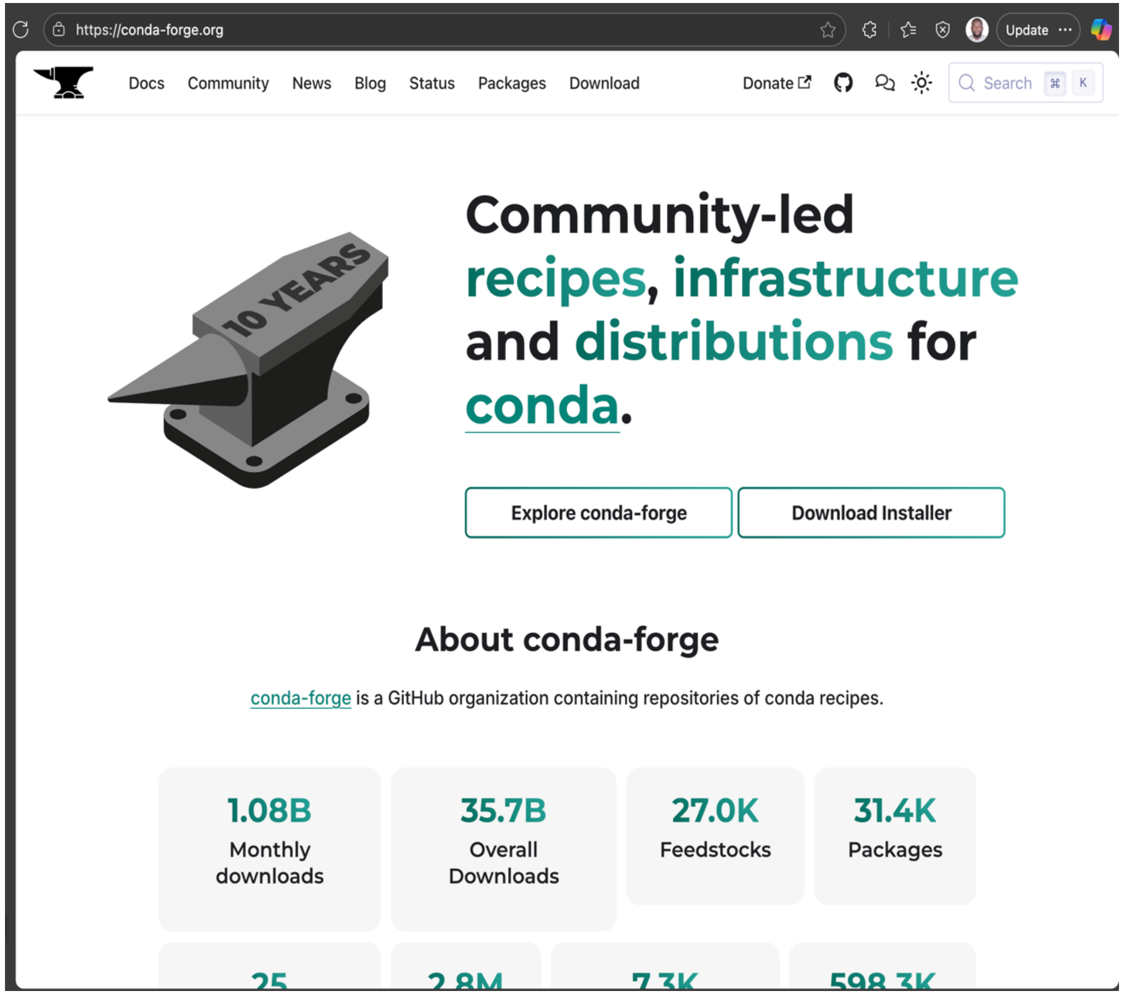

Community-led **recipes**, **infrastructure** and **distributions** for **conda**.

About conda-forge

conda-forge is a GitHub organization containing repositories of conda recipes.

1.08B Monthly downloads

35.7B Overall Downloads

27.0K Feedstocks

31.4K Packages

- Search the web (“conda <package name>”)
- Check conda-forge or Anaconda.
- See our docs for a tutorial to setup a conda environment.
  i. Load the Anaconda/Miniforge module (`module load Anaconda3` or `module load Miniforge3`)
  ii. Create your conda environment (`conda create <envname>`)
  iii. Activate your conda environment (`conda activate <envname>`)
  iv. Install your packages (`conda install <packagename>`)
  v. Your environment is ready to use

Packages for R

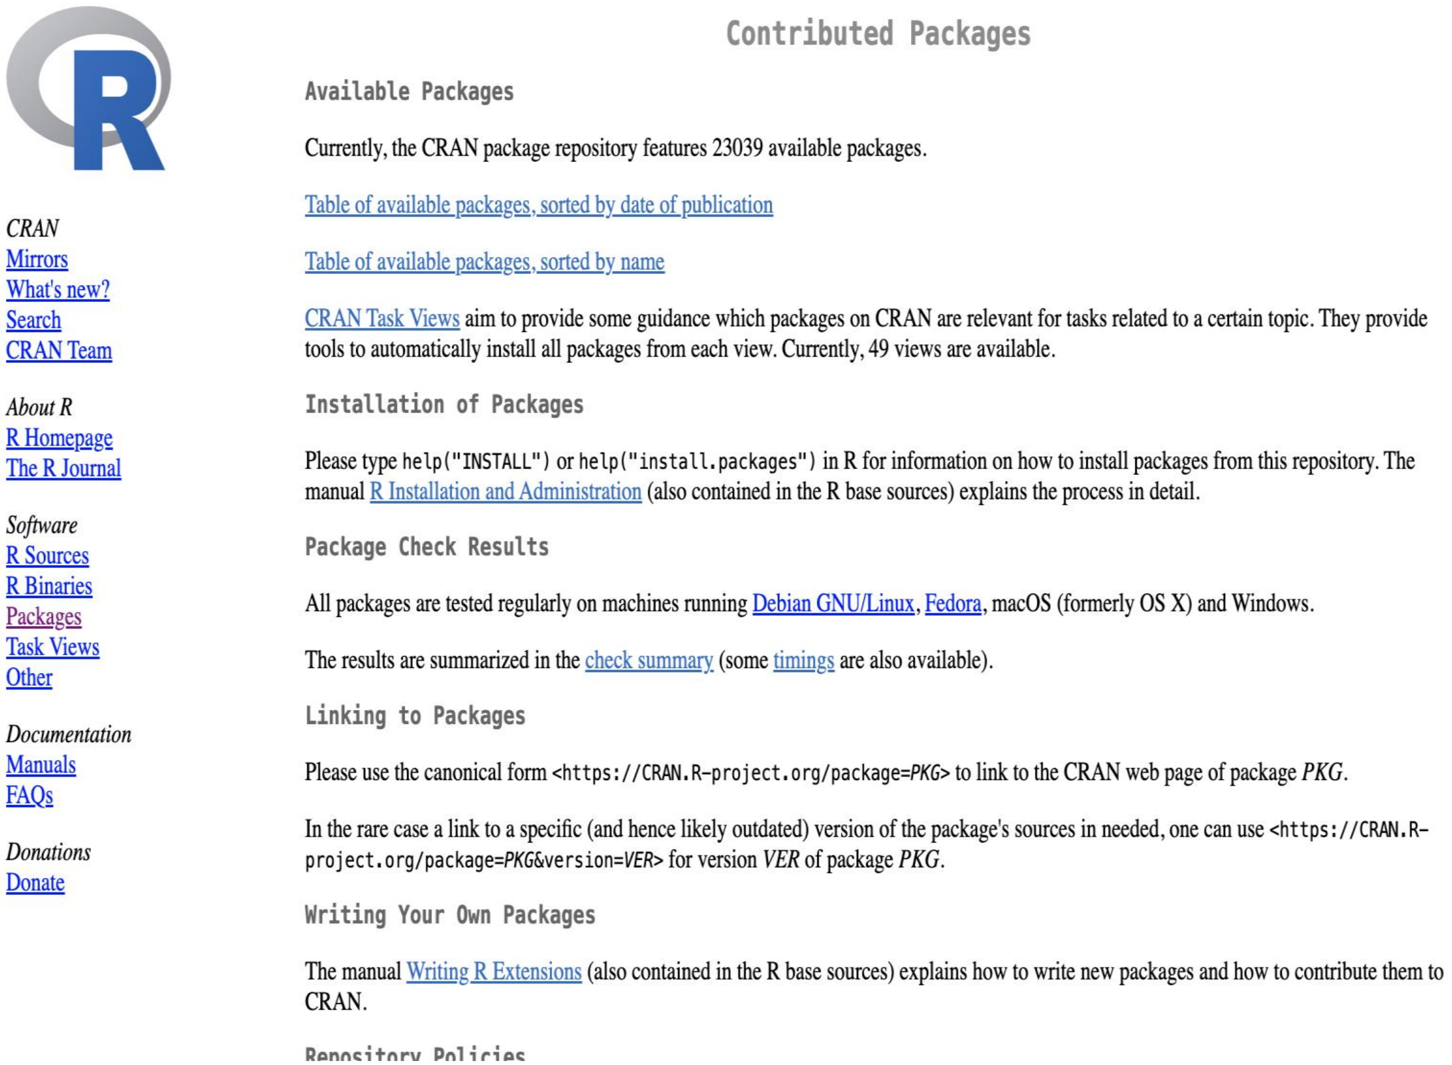

CRAN
Mirrors
What's new?
Search
CRAN Team

About R
R Homepage
The R Journal

Software
R Sources
R Binaries
Packages
Task Views
Other

Documentation
Manuals
FAQs

Donations
Donate

Contributed Packages

Available Packages

Currently, the CRAN package repository features 23039 available packages.

Table of available packages, sorted by date of publication

Table of available packages, sorted by name

CRAN Task Views aim to provide some guidance which packages on CRAN are relevant for tasks related to a certain topic. They provide tools to automatically install all packages from each view. Currently, 49 views are available.

Installation of Packages

Please type `help("INSTALL")` or `help("install.packages")` in R for information on how to install packages from this repository. The manual R Installation and Administration (also contained in the R base sources) explains the process in detail.

Package Check Results

All packages are tested regularly on machines running Debian GNU/Linux, Fedora, macOS (formerly OS X) and Windows.

The results are summarized in the check summary (some timings are also available).

Linking to Packages

Please use the canonical form `<https://CRAN.R-project.org/package=PKG>` to link to the CRAN web page of package *PKG*.

In the rare case a link to a specific (and hence likely outdated) version of the package's sources in needed, one can use `<https://CRAN.R-project.org/package=PKG&version=VER>` for version *VER* of package *PKG*.

Writing Your Own Packages

The manual Writing R Extensions (also contained in the R base sources) explains how to write new packages and how to contribute them to CRAN.

Repository Policies

- Search the web – “<package name> R”
- CRAN - https://cran.r-project.org/
- Bioconductor (Bioinformatics and Genomics packages) - https://www.bioconductor.org
- GitHub
- R Package Managers (`pak, remotes, devtool`)

Running Batch Jobs on Cheaha using Slurm

How do I run jobs on Cheaha?

Jobs

- Interactive Apps (Open On Demand (OOD) Web Apps)
- Batch Jobs (using Slurm)

What is Slurm?

- Slurm – an acronym for Simple Linux Utility Resource Manager

- Slurm is the workload manager or job scheduler we use on Cheaha to submit and manage computing jobs.

- Slurm handles resource allocation, queueing, prioritization, and scheduling on the cluster.

Running Batch Jobs

Running automated, scheduled, unattended (run in absentia) scripts for execution.

Ideal for workflows and pipelines that require more time or computing resources to complete.

Excellent for reproducibility and replication.

Basic Workflow to run a Batch Job on Cheaha

- Stage your data and code onto Cheaha storage.
- Write a Slurm "batch job" script (bash) specifying resources and commands.
- Submit your script with the `sbatch` command.
- Monitor job status with `squeue`.
- After job completion, inspect output/logs, check resource usage, clean up /tmp or scratch if necessary.

Structure of a Slurm Batch Script

```
#!/bin/bash

### Declaring Slurm configuration options and specifying required resources
#SBATCH --job-name=hostname         ### Name of the job
#SBATCH --nodes=1                   ### Number of Nodes
#SBATCH --ntasks=1                  ### Number of Tasks
#SBATCH --cpus-per-task=1           ### Number of Tasks per CPU
#SBATCH --mem=1G                    ### Memory required, 1 gigabyte
#SBATCH --partition=express         ### Cheaha Partition
#SBATCH --time=00:10:00             ### Estimated Time of Completion, 10
minutes
#SBATCH --output=%x_%j.out          ### Slurm Output file, %x is job name, %j
is job id
#SBATCH --error=%x_%j.err           ### Slurm Error file, %x is job name, %j
is job id

### Running the command `hostname`
Hostname
```

Sample Slurm Jobs

- Example 1: A Simple Slurm Batch Job
- Example 2: Sequential Job
- Example 3: Parallel Jobs
- Example 4: Array Jobs
- Example 5: Multithreaded or Multicore Job
- Example 6: GPU Jobs
- Example 7: Multinode Jobs

Live Demo

How can I maximize (efficient) use of Cheaha?

- `seff` is a tool for exploring resource usage efficiency
  - **S**lurm **EFF**iciency
  - wall time: select a partition with a time limit just above this value
  - cpu efficiency: aim for >80%
  - memory efficiency: aim for 80-90%
  - Lower than ~80% for cpu/memory means you are asking for too much, and your jobs are spending longer in the queue than they must.
  - Memory above ~90% is risky, a memory spike could cause your job to be canceled. This can happen depending on your software's programming.
  - Be explicit with resource requests (CPUs, memory, time, partition) helps Slurm schedule efficiently.

Research Computing Services

RC Services and Resources

Compute
Compute nodes for large-scale jobs
(via OOD and Slurm)

Data Transfer Tools
Globus, Rclone, FileZilla, etc.

Storage Offerings
GPFS, LTS

cloud.rc
Based on OpenStack
https://cloud.rc.uab.edu

GitLab
https://code.rc.uab.edu

Software Modules & Containers
Pre-installed software modules and Singularity

Support Requests – Email
support@listserv.uab.edu

Documentation
https://docs.rc.uab.edu

Office Hours
Find Zoom links at:
https://docs.rc.uab.edu/#how-to-contact-us

DSJC Course – Syllabus
Link here

Training Resources
https://docs.rc.uab.edu/education/training_resources/

Storage

Where do I store my data?

		Access Speed	Quota	Use Cases & Notes
Node-local	/local	Fastest	1-6 TB	IO-bound processing like AI
Scratch	/scratch/$USER	Fast	Up to 100 TB	Data being processed in short term
GPFS	/home/$USER /data/user/$USER /data/project/…	Fast	5 TB Personal 25 TB PI/Core	• Data being processed in medium term • Collaborative analysis
LTS	Use Globus, rclone, s5cmd, etc.	Mid	5 TB Personal 75 TB PI/Core	• Raw data • Resource intensive intermediate data • Final products • Data sharing
UAB IT Enterprise Storage		Slow		• Office/Lab Productivity • Managed by UAB Enterprise IT, outside RC • https://www.uab.edu/it/home/tech-solutions/file-storage/storage-options
Archive		Slowest		Let’s talk!

How do I share data?

Collaboration between colleagues at UAB?	Try hosting data in /data/project/ or LTS
Transfer to another institution?	Try Globus
Transfer between GPFS and LTS?	Try Globus, s5cmd, or rclone

- Globus most robust, with GUI
- s5cmd fastest (parallel), least robust, no GUI
- rclone middle ground, no GUI

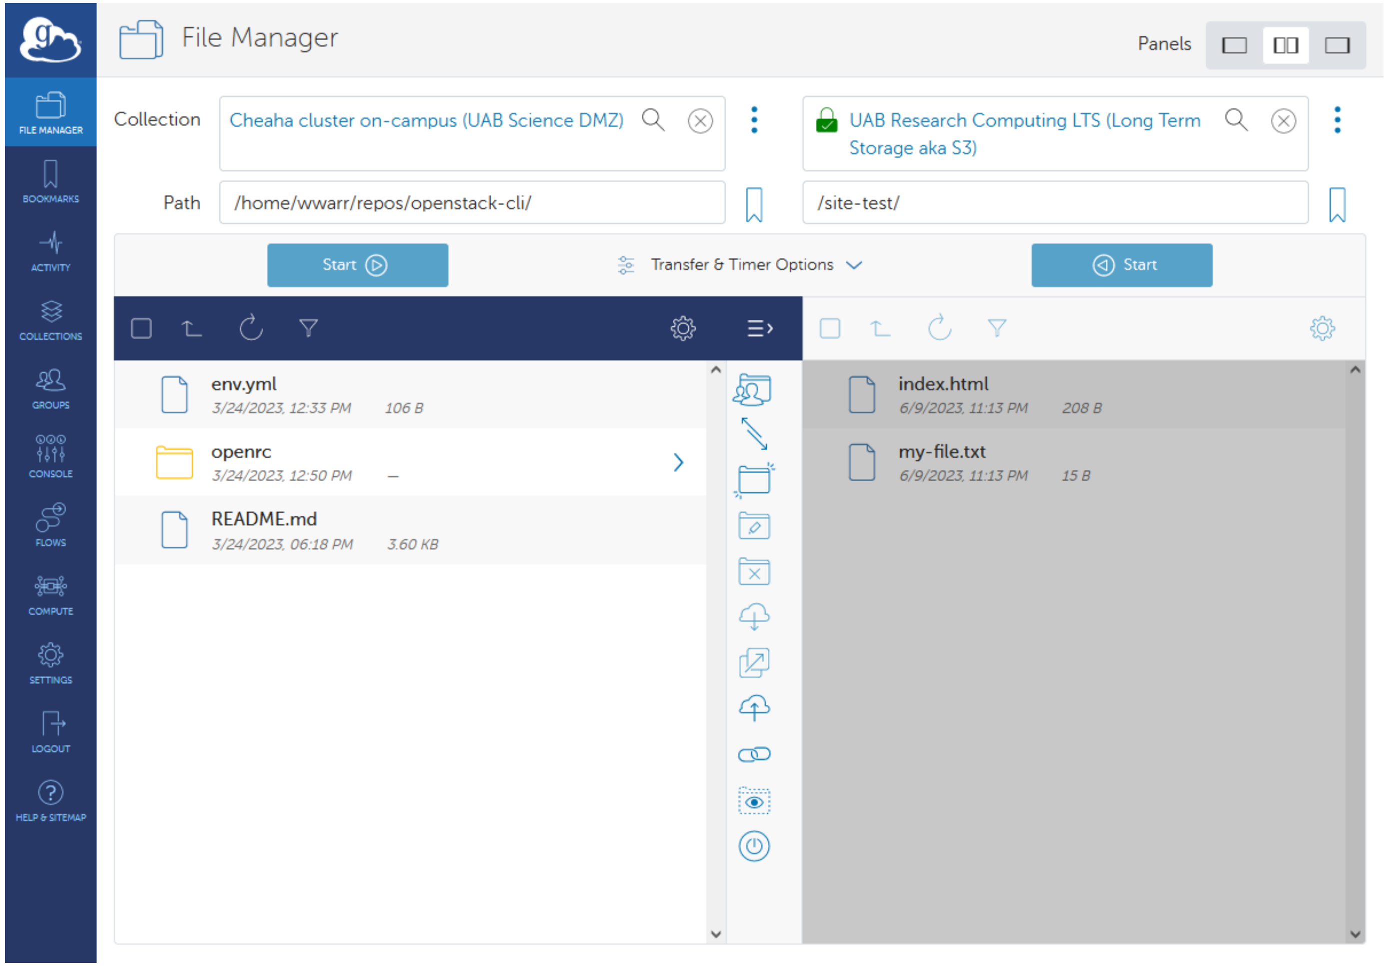

Collection: Cheaha cluster on-campus (UAB Science DMZ)

Path: /home/wwarr/repos/openstack-cli/

Collection: UAB Research Computing LTS (Long Term Storage aka S3)

Path: /site-test/

env.yml	3/24/2023, 12:33 PM	106 B
openrc	3/24/2023, 12:50 PM	–
README.md	3/24/2023, 06:18 PM	3.60 KB

index.html	6/9/2023, 11:13 PM	208 B
my-file.txt	6/9/2023, 11:13 PM	15 B

Good Practices

Good Practices for Using Cheaha

Use the web portal (OOD) for interactive work; use the job scheduler/batch scripts for workflows.

Scientific software and time-consuming commands (more than 5 seconds) should be run only in jobs.

Monitor job usage and efficiency optimize resource requests (seff, sacct).

Use available data management tools. Share data and workflows within your team (project directories).

Version control your scripts and environments for replicability and reproducibility.

Clean up unnecessary, or temp files, and your scratch space regularly.

Reach out for help early: support is available. Email: support@listserv.uab

“Thank you for listening”

Questions?