

# Intro Open OnDemand

**Introduction:** Open OnDemand is a NSF Funded Open Source HPC Project to provide a web-based portal and interface that enables the access and use of High Performance Computing (HPC) systems. ARCC uses Open OnDemand to enable access to Medicine Bow in a user-friendly manner. Discussed in this workshop is each feature of Open OnDemand including the Dashboard, Available Applications, File and Job Management, and Cluster status. There will be some overlap from other modules, but the intent here is to have a one-stop shop for all things Open OnDemand.

## Course Goals:

- How to navigate the Dashboard
- What applications are available and how to use them
- Starting HPC jobs on Open OnDemand

This workshop assumes that participants have already have an understanding of text editors and job scheduling

---

## Sections:

1. [Open OnDemand Dashboard](#)
  2. [Applications in Open OnDemand](#)
  3. [Job Composer](#)
  4. [Interactive Applications](#)
-

# Open OnDemand Dashboard

Discussed in this section of the workshop is getting started with the Open OnDemand interface for MedicineBow, including logging and navigating the dashboard. This is more of an overview of the application and later sections will be the use of the available options found on the Dashboard.

- 
- [Logging In](#)
  - [Overview of the Dashboard](#)
  - [Message of the Day](#)
  - [Pinned Applications](#)
  - [Top Menu - Left Side](#)
  - [Top Menu - Right Side](#)
  - [Next Steps](#)
- 
-

# Logging In

To login to the MedicineBow Open On Demand instance, you will first need to open a web browser (ARCC recommends using Google Chrome over other browsers to limit performance issues) and navigate to <https://medicinebow.arcc.uwyo.edu>.

- You will then be required to enter your ARCC username and password.
- You will also be required to enter your two-factor authentication

You must have an ARCC account before attempting to login

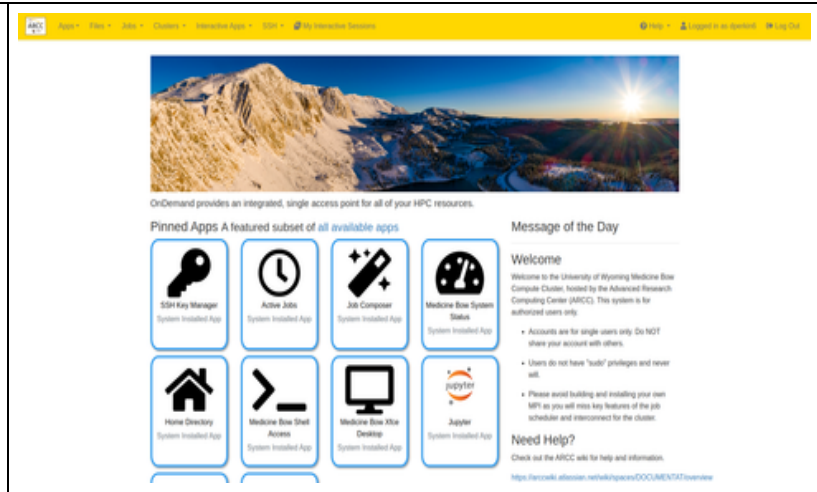


The screenshot shows the ARCC UWyo login interface. At the top is the ARCC logo, which includes a stylized mountain range and the text 'ARCC' and 'UW'. Below the logo is the text 'ARCC UWyo' and 'Sign in using your UWyo credentials'. The login form contains a 'Username' input field, a 'Password' input field with a toggle for visibility, a 'Remember me' checkbox, a blue 'Sign in' button, a 'Forgot' link, and a 'Non-UWyo Users' button.

# Overview of the Dashboard

On the dashboard, you may notice several things:

- The message of the day, that provides ARCC announcements
- A link to the ARCC help pages
- Pinned applications
- A top menu with options for the applications available on Open On Demand
- Information about your login username and an option to log out



## Message of the Day

The message of the day contains important information of users. There are static and some dynamic information. First it tells you what system you are on. Then some policy information about appropriate usage.

Occasionally, there will be an announcement about systems maintenance, or other news that can be important for users to know.

ARCC advises everyone to read the message of the day before diving in to running jobs.

### Message of the Day

#### Welcome

Welcome to the University of Wyoming Medicine Bow Compute Cluster, hosted by the Advanced Research Computing Center (ARCC). This system is for authorized users only.

- Accounts are for single users only. Do NOT share your account with others.
- Users do not have "sudo" privileges and never will.
- Please avoid building and installing your own MPI as you will miss key features of the job scheduler and interconnect for the cluster.

---

# Pinned Applications

The Pinned applications are some of the most commonly used applications, with the basic functions of Open OnDemand of:

- ssh key management
- Job management
- System status
- Files
- Shell Access
- Interactive Desktop

Along with other interactive applications that ARCC has available




---

# Top Menu - Left Side

On the left side of the top menu, you will see drop-down menus for all options available in MedicineBow Open OnDemand beyond what is pinned on the home page.

- ARCC Icon - Back to the dashboard
- Apps - Gives access to some of the options found in pinned apps.
- Files - Access MedicineBow files
- Jobs - Shows your active jobs and lets you create new jobs
- Clusters - shows status of cluster and nodes as well as lets you open a web based terminal
- Interactive Apps
- My Interactive Sessions



## Top Menu - Right Side

<p>On the right side of the top menu, you will see a drop-down menu for help, both the ARCC help pages and the ARCC service portal to send in a support ticket.</p> <p>The other sections of the right side of the top menu shows you who you are logged in as and an option to log out of the web application.</p>	<p>Help ▾    Logged in as dperkin6    Log Out</p> <p>Help ▾</p> <ul style="list-style-type: none"><li>Restart Web Server</li><li>Documentation &amp; Support</li><li>ARCC Wiki</li><li>Support</li></ul>
---	--

# Applications in Open OnDemand

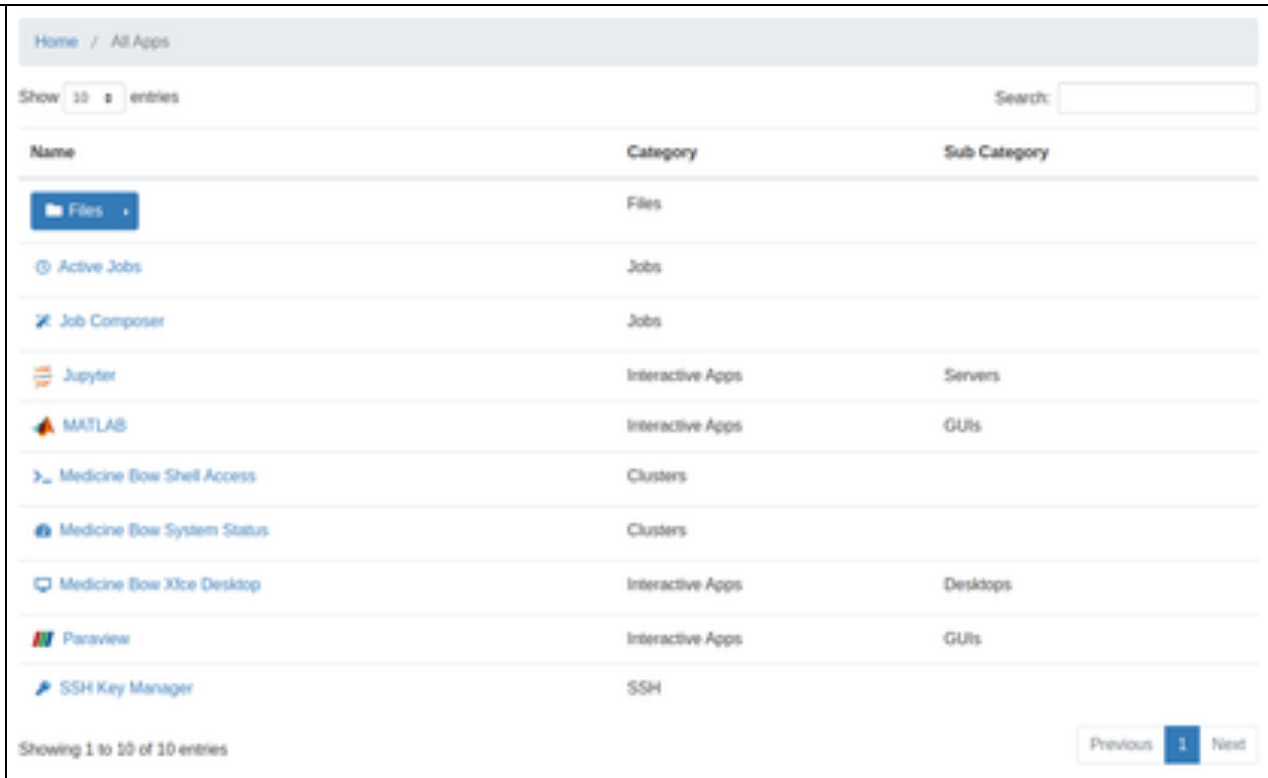
Open OnDemand makes it easy to access your favorite apps for data visualization, simulations, modeling, and more. Discussed in this module is where to find the available applications on MedicineBow Open OnDemand and what the purpose of each is. Some apps are covered more extensively in other tutorials and will be referenced where appropriate.

- 
- [All Applications](#)
  - [The Files Category and App](#)
  - [The Jobs Category](#)
  - [Clusters Category](#)
  - [Interactive Apps Category](#)
  - [Next Steps](#)
-

# All Applications

In the drop down menu for “Apps” you can find all available tools you can use within the Open OnDemand interface, by selecting the “All Apps” option. When on this page, you will see icons and names for all available applications, the category they fall under and the subcategory such as:

- Files
- Jobs
- Interactive Apps
  - Servers, GUIs, and Desktops
- Clusters
- SSH



The screenshot shows the 'All Apps' page in the Open OnDemand interface. At the top, there is a breadcrumb 'Home / All Apps' and a search bar. Below the search bar, there is a 'Show 10 entries' dropdown and a search input field. The main content is a table with three columns: 'Name', 'Category', and 'Sub Category'. The table lists 10 applications: Files, Active Jobs, Job Composer, Jupyter, MATLAB, Medicine Bow Shell Access, Medicine Bow System Status, Medicine Bow Xfce Desktop, Paraview, and SSH Key Manager. At the bottom, there is a pagination bar showing 'Showing 1 to 10 of 10 entries' and 'Previous 1 Next'.

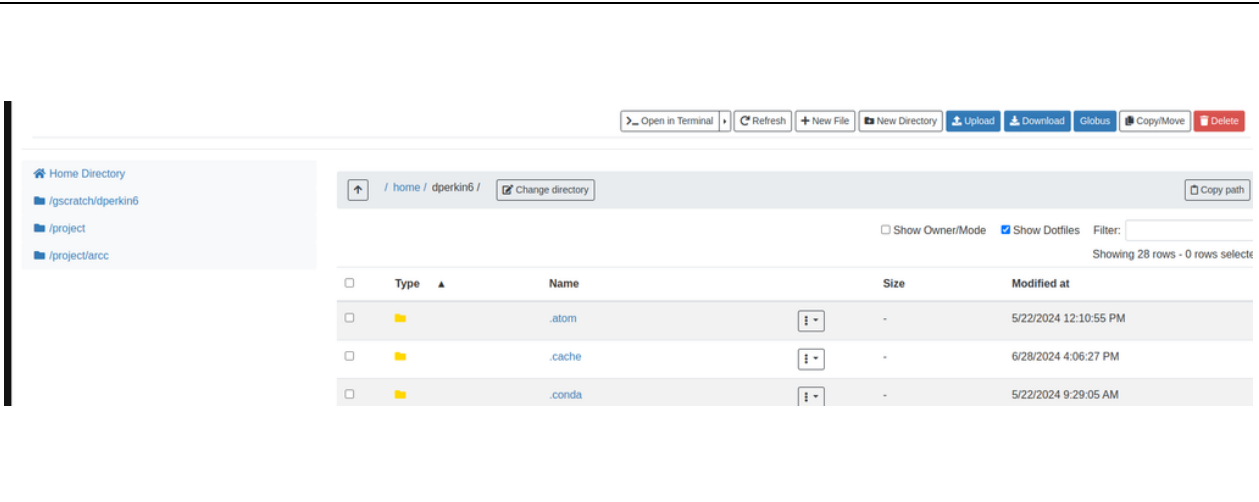
Name	Category	Sub Category
Files	Files	
Active Jobs	Jobs	
Job Composer	Jobs	
Jupyter	Interactive Apps	Servers
MATLAB	Interactive Apps	GUIs
Medicine Bow Shell Access	Clusters	
Medicine Bow System Status	Clusters	
Medicine Bow Xfce Desktop	Interactive Apps	Desktops
Paraview	Interactive Apps	GUIs
SSH Key Manager	SSH	



# The Files Category and App

The files app allows users to view, upload, download, and edit files on the MedicineBow system. When clicking this app, notice that it opens a new browser tab that allows you to close the tab when done without closing the whole application.

The [Intro to Data Transfer](#) workshop provides more details on how to use this application to move data around.



# The Jobs Category

Under Jobs there are two different applications. Active Jobs and the Job Composer. The Active Jobs app shows the status of jobs and many of the details. You can see all jobs on the system or just yours.

The job composer app will be covered more extensively in a later module of this tutorial.

The screenshot shows the MedicineBow Jobs app interface. At the top, there is a table with columns: ID, Name, User, Account, Time Used, Queue, Status, Cluster, and Actions. The first row shows a job with ID 1030454, Name sbatch, User dperkin6, Account arcc, Time Used 00:00:00, Queue inv-arcc, Status COMPLETED, and Cluster Medicine Bow.

ID	Name	User	Account	Time Used	Queue	Status	Cluster	Actions
1030454	sbatch	dperkin6	arcc	00:00:00	inv-arcc	COMPLETED	Medicine Bow	

Below the table, there is a section for job details for sbatch 1030454. The details are as follows:

Cluster	Medicine Bow
Job Id	1030454
Job Name	sbatch
User	dperkin6
Account	arcc
Partition	inv-arcc
State	COMPLETED
Reason	None
Total Nodes	1
Node List	mbcpu-001
Total CPUs	1
Time Limit	1:00
Time Used	0:00
Start Time	2024-07-26 11:17:15
End Time	2024-07-26 11:17:15
Memory	1000M

# Clusters Category

This category opens up a terminal in your web browser and enables users the ability to use it just as if they had a native application on their computer.

It also has a section for system status



```
Host: medicinebow.arcc.uwoy.edu

Medicine Bow

----- Need Help? -----
Check out the ARCC Wiki for more information at https://arccwiki.atlassian.net/
Alternatively, you can e-mail arcc-help@uwoy.edu with any questions.

UM ARCC has office hours-weekly over zoom at https://bit.ly/3BgoGf5
Office hours are 11am-1pm, every Tuesday and 12-2pm Wednesdays. If you need
help but can't make it, email us @ the address above.

See more at https://bit.ly/42MGHNW
-----

>>> Important Messages <<<<

No important messages.

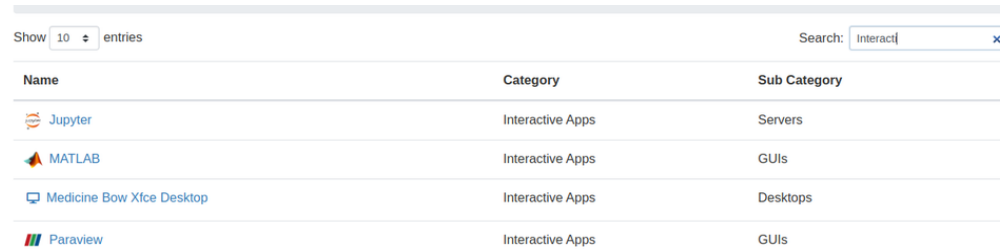
Last login: Thu Jun 27 15:58:50 2024 from 10.83.16.174
+-----+
| arccquota | Block |
+-----+
| Path | Used | Limit | % |
+-----+
| /home/dperkin6 | 977.33 MB | 50.00 GB | 01.91 |
| /gscratch/dperkin6 | 00.00 GB | 05.00 TB | 00.00 |
| /project/arcc | 50.79 GB | 100.00 TB | 00.05 |
+-----+





[dperkin6@mblog2 ~]$
```

# Interactive Apps Category

This category has three sub categories:

- Servers - refers to JupyterLab server covered extensively in another module
- GUIs - Graphical user interfaces are applications that would typically work on a desktop, but on MedicineBow
  - MatLab
  - Paraview
- Desktops - a desktop environment on the cluster



Name	Category	Sub Category
 Jupyter	Interactive Apps	Servers
 MATLAB	Interactive Apps	GUIs
 Medicine Bow Xfce Desktop	Interactive Apps	Desktops
 Paraview	Interactive Apps	GUIs

## Open OnDemand Job Composer

The Open OnDemand Job Composer enables users to create job scripts and submit them without ever using the command line. This is particularly helpful for new users who have never opened a terminal or are unfamiliar with Linux command line behavior. However, some knowledge of Slurm and ARCC's required Slurm directives is needed to successfully submit jobs on MedicineBow. Discussed in this section of the workshop is the use of the Job Composer and common issues that arise with it.

- [First Look at the Job Composer](#)
- [Create a Job With a Template](#)

- [Job Editor](#)
- [Job Options](#)
- [Job Submission](#)
- [In Conclusion](#)
- [Finally The End](#)

---

## First Look at the Job Composer

When opening the Job Composer for the first time, there are a few pop up prompts to follow to get started creating a job. There are also two items in the top menu 1) for Job management and 2) for job templates. Notice that the “New Job” button allows user to select a previously created job with the path, or from a Default Template as well as a new one if there is a job in the selections below.



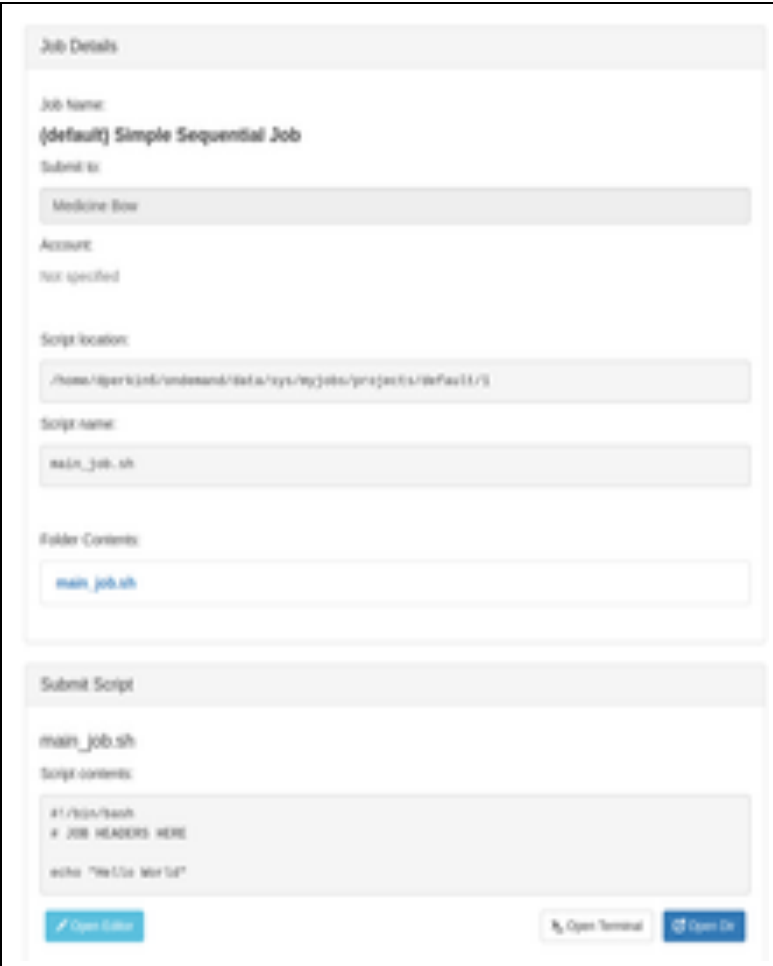
# Create a Job With a Template

If someone has never created a batch job in Slurm before, a template is a good place to get started into seeing how one would look like. In this example, we will create a Default Template and then edit it in the steps below.

After clicking the “New Job” Button and selecting the “Default Template” option, we are presented with a new box on the right of the first look. Notice the parameters already set:

- Job Name
- Account
- Along with script location and name
- As well as a preview of what the job looks like

This default will not run on ARCC systems “as-is”. Edits will be required.



The screenshot displays the Slurm Job Editor interface. It is divided into two main sections: "Job Details" and "Submit Script".

**Job Details:**

- Job Name:** (default) Simple Sequential Job
- Submit to:** Medicine Box
- Account:** Not specified
- Script location:** /home/igperkins/understand/submit/xyz/myjobs/projects/Defaul11/2
- Script name:** main\_job.sh
- Folder Contents:** main\_job.sh

**Submit Script:**

main\_job.sh

Script contents:

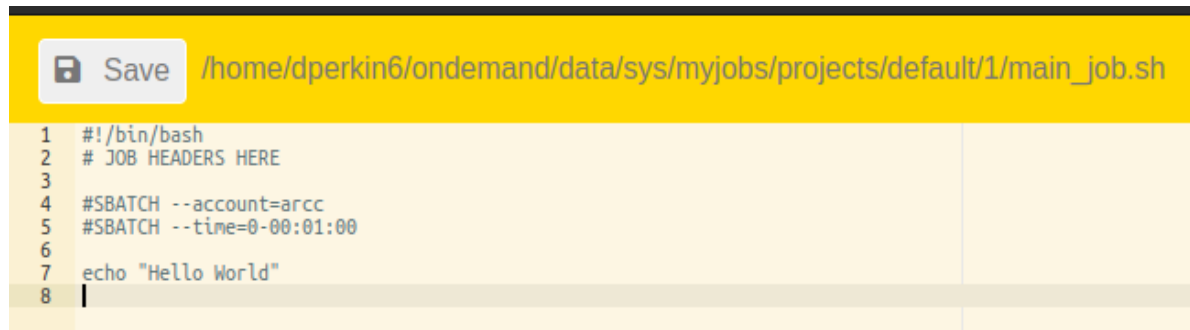
```
#!/bin/bash
# JOB HEADERS HERE

echo "Hello World"
```

At the bottom of the "Submit Script" section, there are three buttons: "Open Editor" (with a pencil icon), "Open Terminal" (with a terminal icon), and "Open On" (with a play icon).

## Job Editor

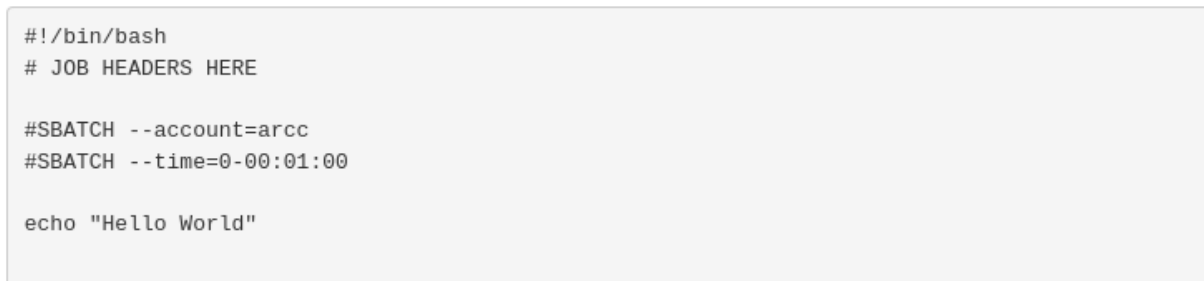
When we click the light blue button on the bottom of the right box, we get a new window to pop up with a text editor-like appearance. If we then look at ARCC's documentation on [Required Inputs](#), we know what to add to get a job to run by adding the account and time Slurm directives.



The screenshot shows a text editor window with a yellow header bar. The header bar contains a 'Save' button and the file path: `/home/dperkin6/ondemand/data/sys/myjobs/projects/default/1/main_job.sh`. The main area of the window displays a shell script with the following content:

```
1 #!/bin/bash
2 # JOB HEADERS HERE
3
4 #SBATCH --account=arcc
5 #SBATCH --time=0-00:01:00
6
7 echo "Hello World"
8 |
```

If we click save and go back to the Job Composer page, we should see our changes in the preview.



The screenshot shows a preview window with a light gray background. It displays the same shell script content as the previous screenshot:

```
#!/bin/bash
# JOB HEADERS HERE


#SBATCH --account=arcc
#SBATCH --time=0-00:01:00

echo "Hello World"
```

---

## Job Options

Now that our job will run, we still might need to make some changes to the job options that will work best for us. If we click the “Job Options”



button on the left box we get presented with a new view:

- We can rename the job to “Hello\_world\_Template” that makes more sense for us to find later
- Leave cluster and job script fields as the default because that is the system we want and the file we just edited
- We can add account here instead of in the script, but it’s just a good practice to have it in the script

### Job Options

Name

Cluster

Specify job script

Files larger than 65KB are omitted for the job script field

Account

Account is an optional field. If not set, the account may be auto-set by the submit filter.

Job array specification

Job arrays are optional. e.g. 1-10

Copy environment

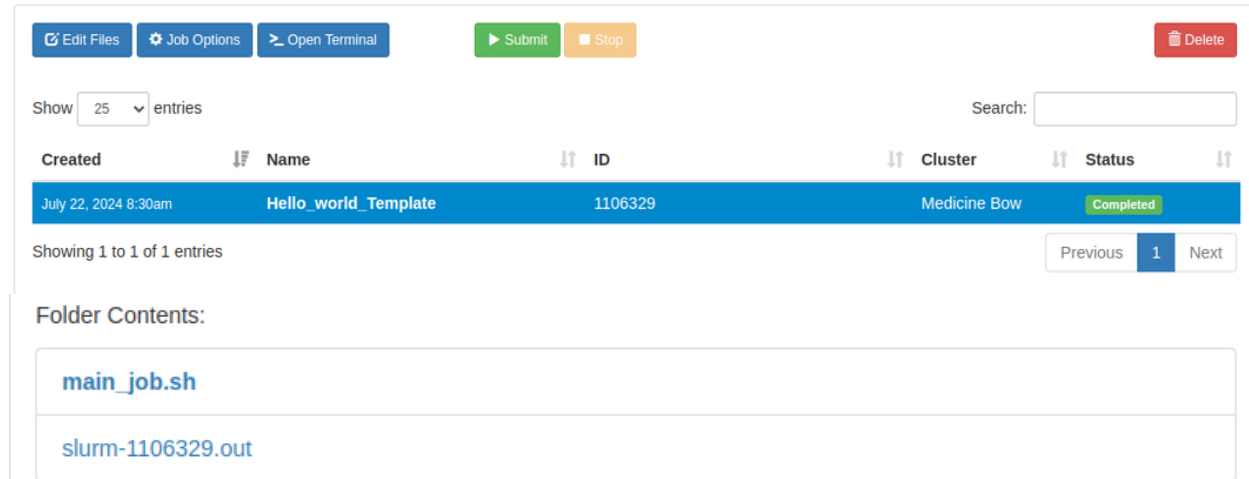
## Job Submission

With all of the edits made, we can now submit the job to the cluster to run. The steps to do so are:

1. Make sure the Job is selected in the left box
2. Click the green “Submit” button



If we refresh the page, we see in the left box that this simple job has been completed and a file for our output was created on the right box



The screenshot shows the Job Composer interface. At the top, there are buttons for 'Edit Files', 'Job Options', 'Open Terminal', 'Submit', 'Stop', and 'Delete'. Below these is a search bar and a 'Show 25 entries' dropdown. A table lists job entries with columns for 'Created', 'Name', 'ID', 'Cluster', and 'Status'. One entry is highlighted in blue: 'Hello\_world\_Template' with ID '1106329' on the 'Medicine Bow' cluster, with a 'Completed' status. Below the table, it says 'Showing 1 to 1 of 1 entries' and has 'Previous', '1', and 'Next' navigation buttons. Under 'Folder Contents:', there are two files listed: 'main\_job.sh' and 'slurm-1106329.out'.

Created	Name	ID	Cluster	Status
July 22, 2024 8:30am	Hello_world_Template	1106329	Medicine Bow	Completed

## In Conclusion

This module was a simple tutorial for using the Job Composer in Open OnDemand. This tool helps to send jobs to the MedicineBow cluster without the need for using a terminal. That said, the concepts for job submission with Slurm still apply to this tool and it is important to understand what each are doing, but most importantly:

- ❑ ARCC’s required inputs
- ❑ Other Slurm directives that are important to the success of the job
- ❑ How to use Lmod to load software
- ❑ Awareness of where the locations on the system where the script is created and output are placed



If used wisely, the Open OnDemand Job Composer can be an easy-to-use tool for creating jobs on MedicineBow and the other tutorials on Intro to Linux, Text Editors, Lmod, and Job Submission are extremely beneficial for the successful use of the tool.

---

## Interactive Applications

In the applications module in this tutorial the interactive applications were mentioned as available, but there are some nuances to getting them started and working for your use case. On this module we will go over getting the interactive desktop, MATLAB, and Paraview interactive applications working. Jupyter Notebooks will be covered in another workshop.

---

- [Interactive Desktop](#)
  - [Configuring the Desktop](#)
  - [Launching the Desktop](#)
  - [Exploring the Desktop](#)
  - [MATLAB](#)
  - [Parallel Computing in MATLAB](#)
  - [ParaView](#)
  - [Using Paraview](#)
  - [Managing Interactive Sessions](#)
  - [Next Steps](#)
- 

## Interactive Desktop

Open OnDemand's interactive desktop application enables users to use the MedicineBow HPC system as if it were a regular desktop computer that many are used to rather than a command line interface (CLI). However, since MedicineBow is a Linux based system, so

is the interactive desktop and may not be as familiar as a Windows or Mac desktop. On MedicineBow, the Xfce desktop environment is used to facilitate this functionality.

Before you use the interactive desktop be prepared that it will prompt you for:

- Project/Account
- Number of hours
- Desktop Configuration for your CPU and RAM requirements
- GPU requirements

---

## **Configuring the Desktop**

In this simple example we configure the desktop with:

- account
- 1 hour
- 1 cpu, 4GB
  - There are more options in the dropdown
- No GPU
  - in the drop down the options are A30, L40s, and H100

## Medicine Bow Xfce Desktop

This app will launch an interactive desktop on one or more compute nodes. You will have full access to the resources these nodes provide. This is analogous to an interactive batch job.

Project/Account

arcc

Pick the slurm account you want to use.

Number of hours

1

The maximum number of hours your session will run.

Desktop Configuration

1 CPU, 4GB memory

Select the CPU and memory configuration for the session.

GPU Type

None - No GPU

GPU type to allocate to this session. Requesting a GPU may cause the job to queue and not start immediately.

Launch

\* The Medicine Bow Xfce Desktop session data for this session can be accessed under the [data root directory](#).

## Launching the Desktop

When launching the desktop, it will not open immediately. What is happening is that the system is allocating the hardware needed and then submitting a job to a compute node to run.

Once it is ready, you will see a green bar that says it's running.

To actually launch the desktop, you have to click the Launch button on the bottom of the box

Launch Medicine Bow Xfce Desktop

Medicine Bow Xfce Desktop (1106332)

1 node | 1 core | Running

Host: mbcpu-001

Cancel

Created at: 2024-07-22 11:25:06 MDT

Time Remaining: 58 minutes

Session ID: 3d72fce8-8f36-4481-91c0-2354e244df75

Compression

0 (low) to 9 (high)

Image Quality

0 (low) to 9 (high)

Launch Medicine Bow Xfce Desktop

View Only (Share-able Link)

## Exploring the Desktop

Once launched, the desktop appears in the browser window, with familiar apps in the bottom menu bar like:

File browser




Internet



Terminal



Some GUI applications can be launched in this desktop or manage files, etc.



## MATLAB

Launching MATLAB is a similar process, the options are a little different where version of the software is one of them and the CPU and RAM options can be manually selected instead of a drop down.

## MATLAB

This app will launch a [MATLAB](#) GUI on the Medicine Bow cluster. You will be able to interact with the MATLAB GUI through a VNC session.

### Matlab Version

Select which version of Matlab you would like to load.

### Project/Account

Pick the slurm account you want to use.

### Number of hours

The maximum number of hours your session will run. (1-100)

### Number of CPUs

Number of CPUs to allocate to this session. (1-96)

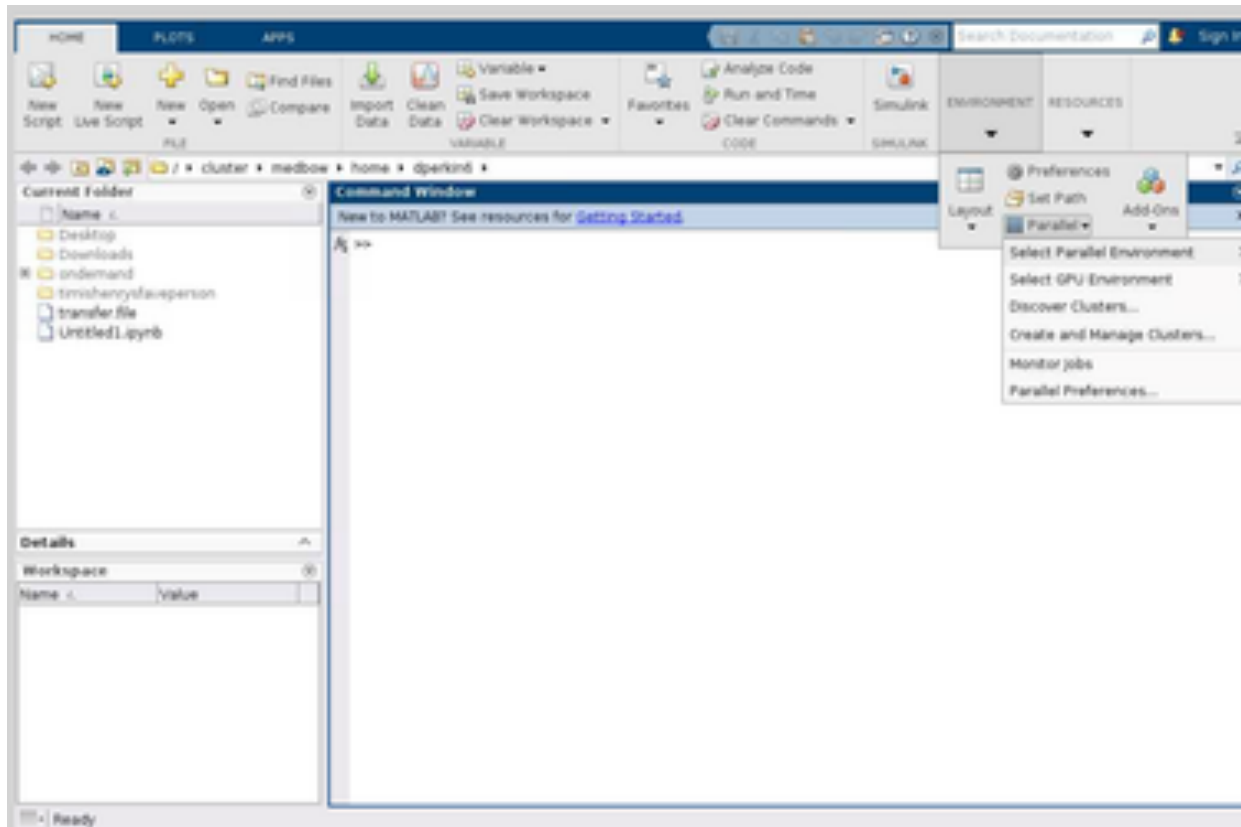
### Total Memory (GB)

Total amount of memory in GB to allocate to this session. (8-512)

### GPU Type

# Parallel Computing in MATLAB

If you've used Matlab before, you'll be familiar with the window you see after launching. However, In order to use the parallel functionality of MATLAB, you have to set your environment.



Also, be aware if you are running parallel, that by default, Matlab restricts you to 12 worker threads, so you'd need to override this if you needed by setting the `poolobj = parpool('local', 24);`

---

# ParaView

Once again the process of launching the Paraview GUI is similar, but with different options:



## Paraview

This app will launch a [ParaView](#) GUI on the Medicine Bow cluster using a **shared node**. You will be able to interact with the ParaView GUI through a VNC session.

Paraview Version

This defines the version of Paraview you want to load.

Project/Account

Pick the slurm account you want to use.

Max Session Duration (hours)

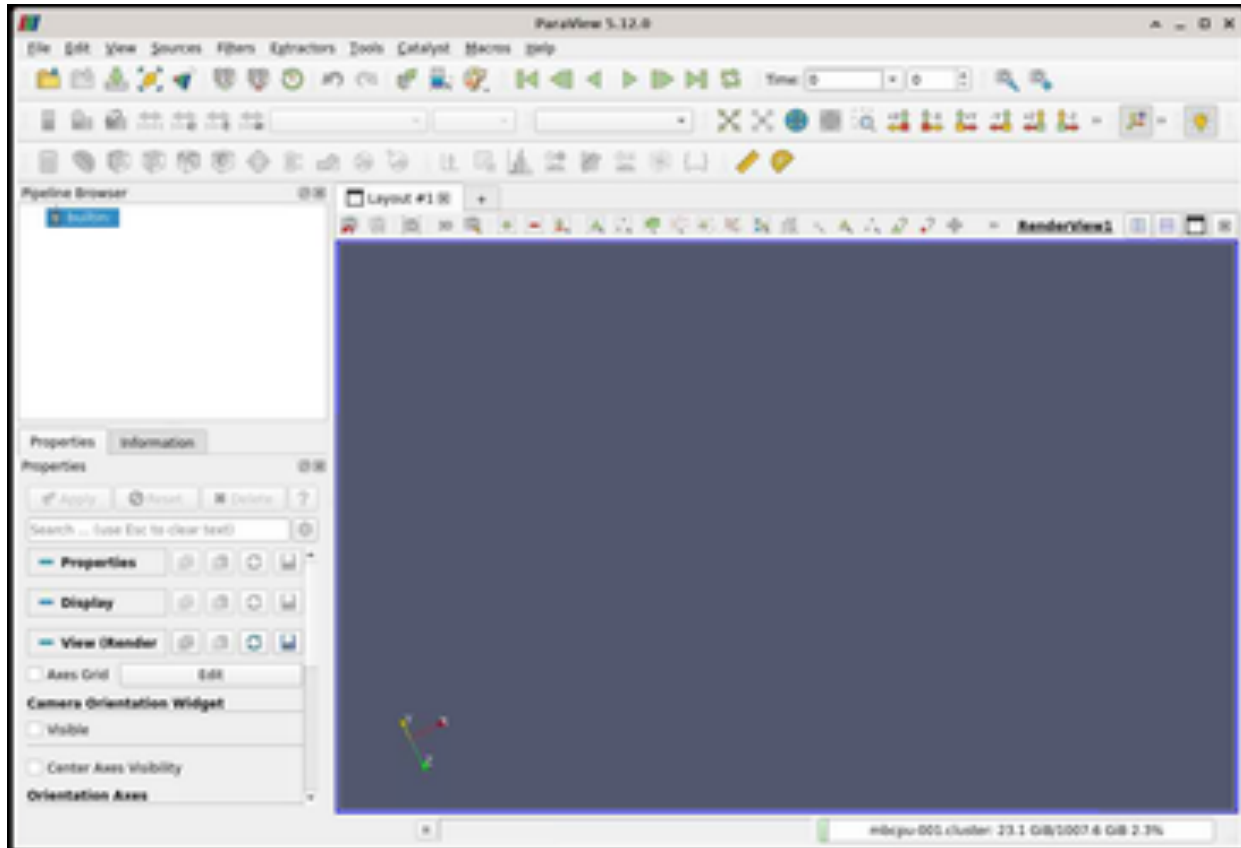
The maximum number of hours your session can run (1-168)

\* The Paraview session data for this session can be accessed under the [data root directory](#).

---

## Using Paraview

Once the application is launched, you will see the application running in your browser window.



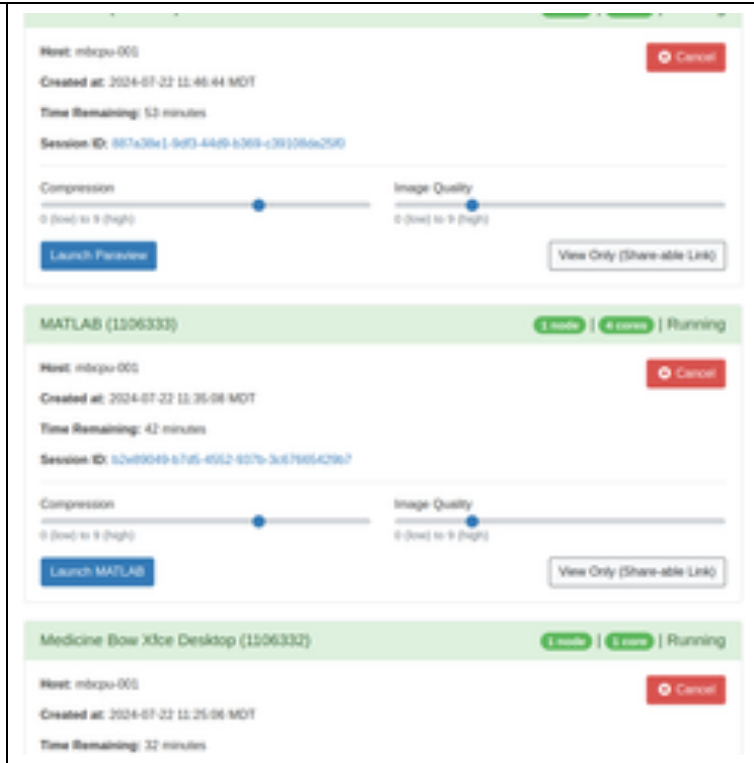
# Managing Interactive Sessions

It is a good practice to keep the number of interactive sessions you are working on in check. As each of these are using resources on MedicineBow, if you have many open sessions that are idle, you will be preventing other users from using them.

You can free up those resources by cancelling your sessions using the



“Cancel” button



## Next Steps

Previous

[Open OnDemand Job Composer](#)

Workshop Home

[Intro to Open OnDemand](#)