

OPEN ONDEMAND 101

Ryan DeRue, Senior Computational Scientist

000 101

Outline

What to Expect from Open OnDemand 101?

Objectives

- Understand how science gateways like Open OnDemand address the challenges faced in typical high-performance computing (HPC) work
- Ensure you know how to login to the gateway associated with your HPC resource here at Purdue
- Become comfortable with the purpose and location of different features in Open OnDemand
- Feel confident running your own workflows in Open OnDemand


000 101

Lowering the Barrier to HPC

Lowering the Barrier to HPC

Layout of a Cluster


- HPC systems are often organized into clusters
 - We call them clusters because they are a multitude of independent nodes that can communicate with one another
 - Because each node is independent, we can have different types of nodes with specialized hardware
- When you login to a cluster, typically you will land on login node along with all the other users on the cluster
 - In order to carve out dedicated resources for your work, you need to ask the scheduler to allocate you resources on the compute nodes



Lowering the Barrier to HPC

HPC Presents a Steep Learning Curve

- Working within HPC in this traditional manner presents many challenges to new users
 - Typically use an SSH client/Linux terminal
 - Learning new commands for interacting with the scheduler
 - Software runs without a graphical user interface (GUI)
- This can lead to new domain scientists spending most of their time learning how to use the system as opposed to doing science




Lowering the Barrier to HPC

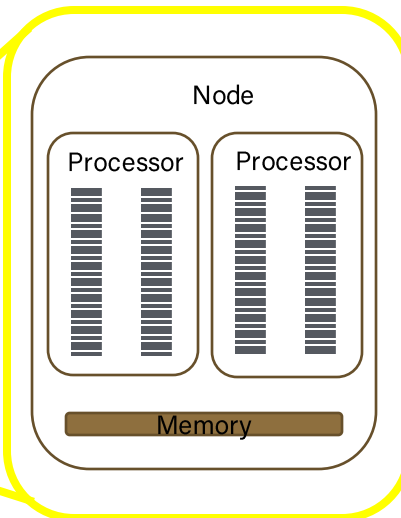

Science Gateways

- To make these resources easier to use, there has been a push for the creation of “science gateways”
- Science Gateways are typically web based interfaces that allow users to run science/engineering applications on HPC resources from their web browser
- Open OnDemand is an example of such a science gateway

Your Computer



Scheduler



Compute Nodes

000 101

Using Open OnDemand: Dashboard

Using Open OnDemand: Dashboard

Logging into Open OnDemand

- We offer Open OnDemand on all Purdue community clusters, however each cluster has its own access gateway
 - Each of these can be found by navigating to: `gateway.<cluster>.rcac.purdue.edu`
 - For example, for the Gilbreth cluster,
- Upon navigating to the gateway for your cluster, you will need to login using BoilerKey
- Once you are logged in, you will land on your dashboard page
 - My dashboard for `gateway.gilbreth.rcac.purdue.edu` is on the right

Gilbreth - Gateway Files Jobs Interactive Apps Documentation Cluster My Interactive Sessions Develop Help Logged in as rderue Log Out

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

See documentation for more help in using OnDemand.

Current Quota Usage

/home/rderue Using 2.94 GB of quota 25 GB 11%	No file count limit.	Updated 7 months, 1 week, 4 days, 21 hours, and 19 minutes ago
/scratch/gilbreth/rderue Using 0 Bytes of quota 1 KB	No file count limit.	No data yet
/depot/lslipche Using 994 GB of quota 2 TB 48%	No file count limit.	Updated 49 minutes ago

* Please allow up to 15 minutes for these numbers to update.

powered by **OPEN OnDemand**

Dashboard version: v1.32.0

Using Open OnDemand: Dashboard

Getting Familiar with The Dashboard

1. Current Quota Usage

- Shows the remaining space in each of your different filesystem quotas
- Includes your file count quota as well!
- To learn more about storage quotas see our user guide page on quotas

The screenshot shows the Open OnDemand dashboard interface. At the top, there is a navigation bar with the following items: "Gilbreth - Gateway", "Files", "Jobs", "Interactive Apps", "Documentation", "Cluster", "My Interactive Sessions", "Develop", "Help", "Logged in as rderue", and "Log Out". Below the navigation bar, a message states: "OnDemand provides an integrated, single access point for all of your HPC resources." followed by a link to "See documentation for more help in using OnDemand." The main content area is titled "Current Quota Usage" and is highlighted with a yellow border. It displays three rows of quota information:

Filesystem	Usage	Quota	File Count Limit	Last Updated
/home/rderue	Using 2.94 GB of quota 25 GB 11%	25 GB	No file count limit.	Updated 7 months, 1 week, 4 days, 21 hours, and 19 minutes ago
/scratch/gilbreth/rderue	Using 0 Bytes of quota 1 KB	1 KB	No file count limit.	No data yet
/depot/lslipche	Using 994 GB of quota 2 TB 48%	2 TB	No file count limit.	Updated 49 minutes ago

At the bottom of the dashboard, there is a footer that says "powered by OPEN OnDemand" and "Dashboard version: v1.32.0". A note at the bottom of the quota section reads: "* Please allow up to 15 minutes for these numbers to update."

Using Open OnDemand: Dashboard

Getting Familiar with The Dashboard

2. Directory Shortcuts

- Allows you to quickly navigate between your important directories
- Clicking these will open a file navigator allowing you to view/edit your files or upload/download files to/from the cluster

The screenshot shows the Open OnDemand dashboard interface. At the top, there is a navigation bar with the following items: "Gilbreth - Gateway", "Files", "Jobs", "Interactive Apps", "Documentation", "Cluster", "My Interactive Sessions", "Develop", "Help", "Logged in as rderue", and "Log Out". The "Files" menu is highlighted with a yellow box, showing a dropdown list with three options: "Home Directory", "/scratch/gilbreth/rderue", and "/depot/lslipche". Below the navigation bar, there is a section titled "Current Quota Usage" which displays three rows of data for different directories. Each row includes a progress bar showing the percentage of quota used and a "No file count limit" indicator.

Directory	Quota Usage	File Count Limit	Last Updated
/home/rderue	Using 2.94 GB of quota 25 GB (11%)	No file count limit.	Updated 7 months, 1 week, 6 days, and 17 minutes ago
/scratch/gilbreth/rderue	Using 0 Bytes of quota 1 KB	No file count limit.	No data yet
/depot/lslipche	Using 995 GB of quota 2 TB (48%)	No file count limit.	Updated less than 1 second ago

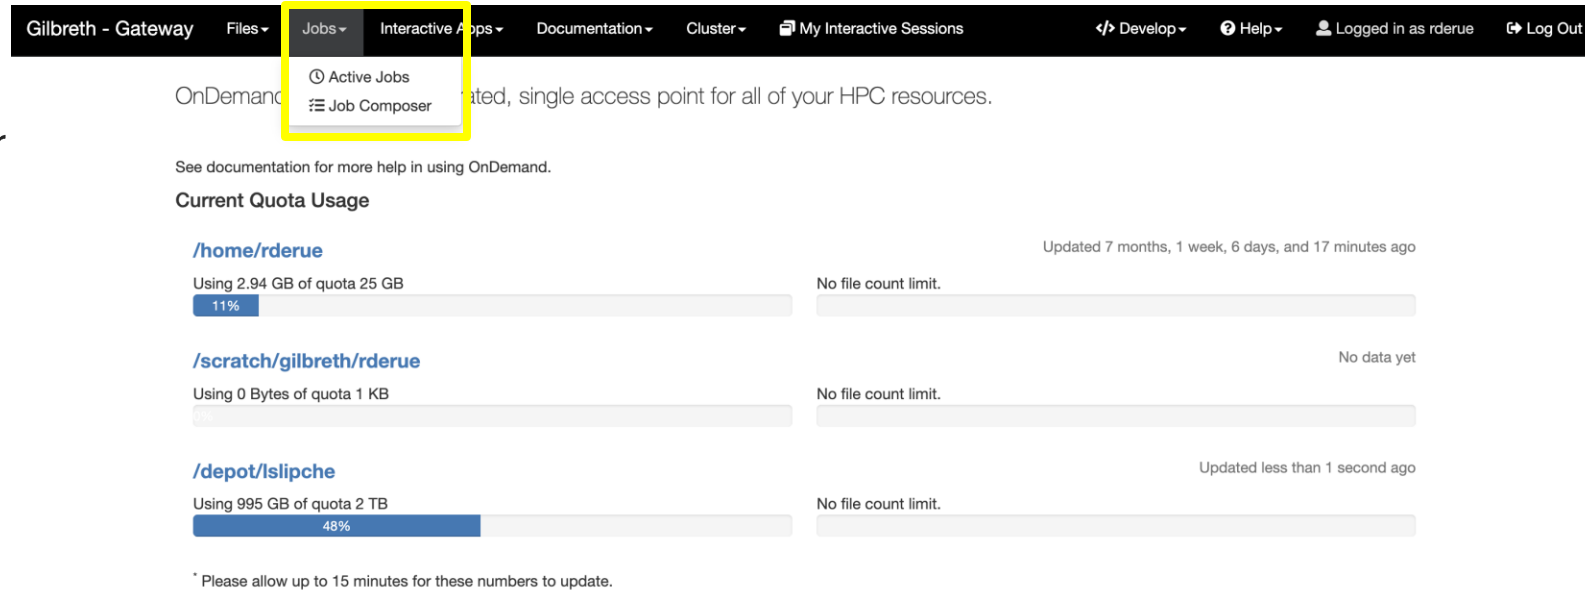
* Please allow up to 15 minutes for these numbers to update.

Using Open OnDemand: Dashboard

Getting Familiar with The Dashboard

3. Jobs Menu

- Provides shortcuts to widgets for creating new Slurm jobs as well as viewing the status of your current ones



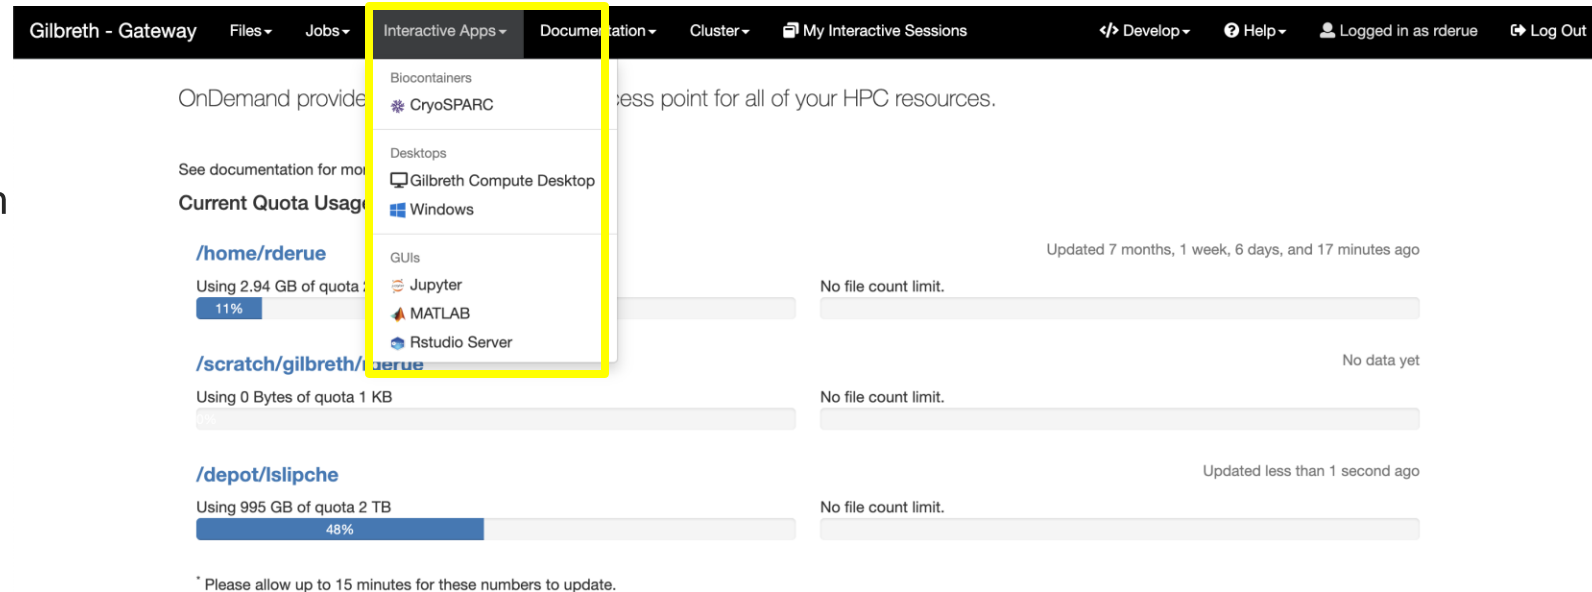
The screenshot shows the Open OnDemand dashboard interface. The top navigation bar includes links for 'Gilbreth - Gateway', 'Files', 'Jobs', 'Interactive Apps', 'Documentation', 'Cluster', 'My Interactive Sessions', 'Develop', 'Help', 'Logged in as rderue', and 'Log Out'. The 'Jobs' menu is highlighted with a yellow box, showing options for 'Active Jobs' and 'Job Composer'. Below the navigation bar, the dashboard displays 'Current Quota Usage' for three paths: '/home/rderue' (11% usage), '/scratch/gilbreth/rderue' (0 Bytes usage), and '/depot/lslipche' (48% usage). Each path includes a progress bar and a 'No file count limit' indicator. A note at the bottom states: '* Please allow up to 15 minutes for these numbers to update.'

Using Open OnDemand: Dashboard

Getting Familiar with The Dashboard

4. Interactive Application Menu

- This dropdown lists all the applications we support on Open OnDemand
- Clicking one of the applications will take you to a form through which you can request the resources you want to use



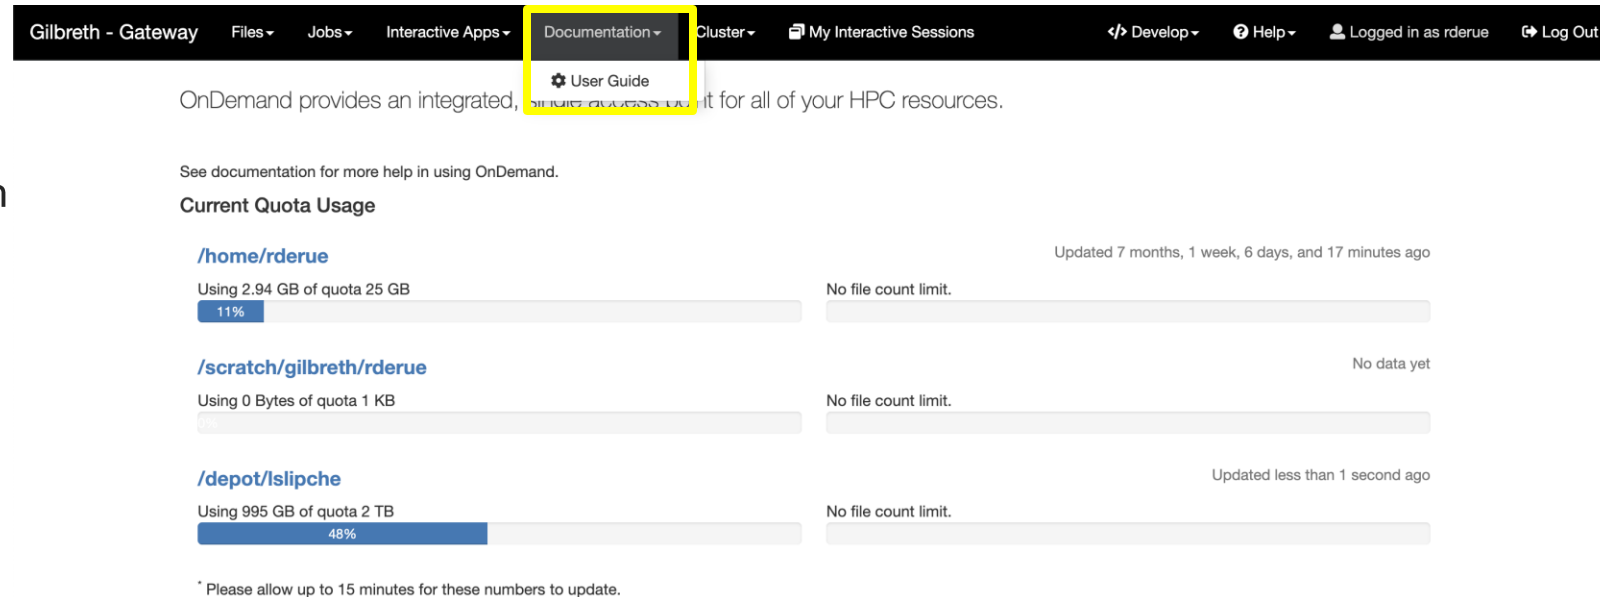
The screenshot shows the Open OnDemand dashboard interface. At the top, there is a navigation bar with the following items: "Gilbreth - Gateway", "Files", "Jobs", "Interactive Apps", "Documentation", "Cluster", "My Interactive Sessions", "Develop", "Help", "Logged in as rderue", and "Log Out". The "Interactive Apps" dropdown menu is open, listing the following applications: "Biocontainers", "CryoSPARC", "Desktops", "Gilbreth Compute Desktop", "Windows", "GUIs", "Jupyter", "MATLAB", and "Rstudio Server". Below the navigation bar, the dashboard displays "Current Quota Usage" for three different paths: "/home/rderue" (Using 2.94 GB of quota, 11%), "/scratch/gilbreth/rderue" (Using 0 Bytes of quota, 1 KB), and "/depot/lsipche" (Using 995 GB of quota, 2 TB, 48%). Each path has a progress bar and a "No file count limit." message. A footer note states: "* Please allow up to 15 minutes for these numbers to update."

Using Open OnDemand: Dashboard

Getting Familiar with The Dashboard

5. Documentation

- Provides a shortcut to our user guide for the cluster whose Open OnDemand instance you are using



The screenshot shows the Open OnDemand dashboard interface. At the top, there is a navigation bar with several menu items: "Gilbreth - Gateway", "Files", "Jobs", "Interactive Apps", "Documentation", "Cluster", "My Interactive Sessions", "Develop", "Help", "Logged in as rderue", and "Log Out". The "Documentation" menu is highlighted with a yellow box, and a sub-menu item "User Guide" is visible below it. Below the navigation bar, the main content area displays the text: "OnDemand provides an integrated, [redacted] for all of your HPC resources." Below this, there is a section titled "Current Quota Usage" with a sub-header "See documentation for more help in using OnDemand." The "Current Quota Usage" section contains three rows of data, each with a path, usage information, and a progress bar. The first row is for "/home/rderue" with 2.94 GB of quota (25 GB total) and 11% usage. The second row is for "/scratch/gilbreth/rderue" with 0 Bytes of quota (1 KB total) and no data yet. The third row is for "/depot/Islipche" with 995 GB of quota (2 TB total) and 48% usage. Each row also indicates "No file count limit." and shows the last update time. A note at the bottom states: "Please allow up to 15 minutes for these numbers to update."

Path	Usage	Quota	File Count Limit	Last Updated
/home/rderue	Using 2.94 GB of quota 25 GB 11%	25 GB	No file count limit.	Updated 7 months, 1 week, 6 days, and 17 minutes ago
/scratch/gilbreth/rderue	Using 0 Bytes of quota 1 KB	1 KB	No file count limit.	No data yet
/depot/Islipche	Using 995 GB of quota 2 TB 48%	2 TB	No file count limit.	Updated less than 1 second ago

Using Open OnDemand: Dashboard

Getting Familiar with The Dashboard

6. Cluster

- Provides options for interacting with the cluster directly
- Gilbreth Shell Access will provide you with a Unix shell on one of the login nodes
- Open XDMoD provides a tool for viewing metrics about the cluster

The screenshot shows the Open OnDemand dashboard for the 'Gilbreth - Gateway'. The top navigation bar includes 'Files', 'Jobs', 'Interactive Apps', 'Documentation', 'Cluster', and 'My Interactive Sessions'. The 'Cluster' menu is highlighted with a yellow box, showing options for 'Gilbreth Shell Access', 'Metrics', and 'Open XDMoD'. Below the navigation, the dashboard displays 'Current Quota Usage' for three paths: /home/rderue (11% usage), /scratch/gilbreth/rderue (0 Bytes usage), and /depot/lslipche (48% usage). Each path has a progress bar and a 'No file count limit' indicator. A note at the bottom states: '* Please allow up to 15 minutes for these numbers to update.'

Using Open OnDemand: Dashboard

Getting Familiar with The Dashboard

7. My Interactive Sessions

- This will take you to a page wherein you can access all your currently running application sessions

Gilbreth - Gateway Files Jobs Interactive Apps Documentation Cluster **My Interactive Sessions** Develop Help Logged in as rderue Log Out

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

See documentation for more help in using OnDemand.

Current Quota Usage

/home/rderue Using 2.94 GB of quota 25 GB 11%	No file count limit.	Updated 7 months, 1 week, 4 days, 21 hours, and 19 minutes ago
/scratch/gilbreth/rderue Using 0 Bytes of quota 1 KB	No file count limit.	No data yet
/depot/lslipche Using 994 GB of quota 2 TB 48%	No file count limit.	Updated 49 minutes ago

* Please allow up to 15 minutes for these numbers to update.

powered by **OPEN OnDemand**

Dashboard version: v1.32.0

000 101

Using Open OnDemand: Files

Using Open OnDemand: Files

A Web-Based File Explorer

1. To access the File Explorer, click on one of the mounted locations listed in the Files dropdown
 - This will open a new tab containing the File Explorer
 - You can also click the blue hyperlink text within the quota widget

The screenshot shows the Open OnDemand web interface. The top navigation bar includes "Gilbreth - Gateway", "Files", "Jobs", "Interactive Apps", "Documentation", "Cluster", "My Interactive Sessions", "Develop", "Help", "Logged in as rderue", and "Log Out". The "Files" dropdown menu is highlighted with a yellow box, showing options for "Home Directory", "/scratch/gilbreth/rderue", and "/depot/lslipche". Below the navigation bar, the "Current Quota Usage" section displays three rows of quota information:

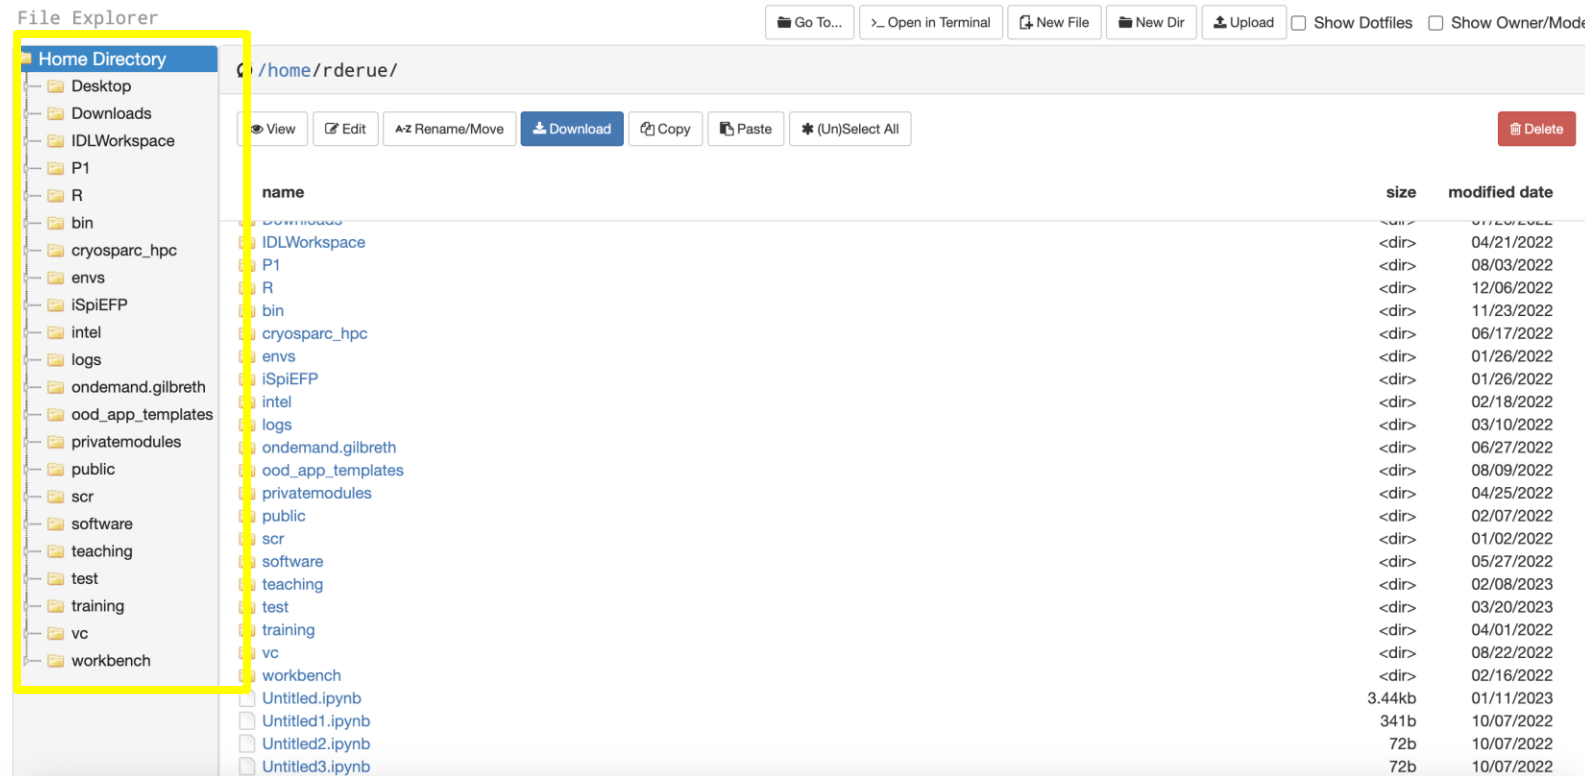
Path	Usage	Quota	File Count Limit	Last Updated
/home/rderue	Using 2.94 GB of quota 25 GB 11%	25 GB	No file count limit.	Updated 7 months, 1 week, 6 days, and 17 minutes ago
/scratch/gilbreth/rderue	Using 0 Bytes of quota 1 KB	1 KB	No file count limit.	No data yet
/depot/lslipche	Using 995 GB of quota 2 TB 48%	2 TB	No file count limit.	Updated less than 1 second ago

* Please allow up to 15 minutes for these numbers to update.

Using Open OnDemand: Files

A Web-Based File Explorer

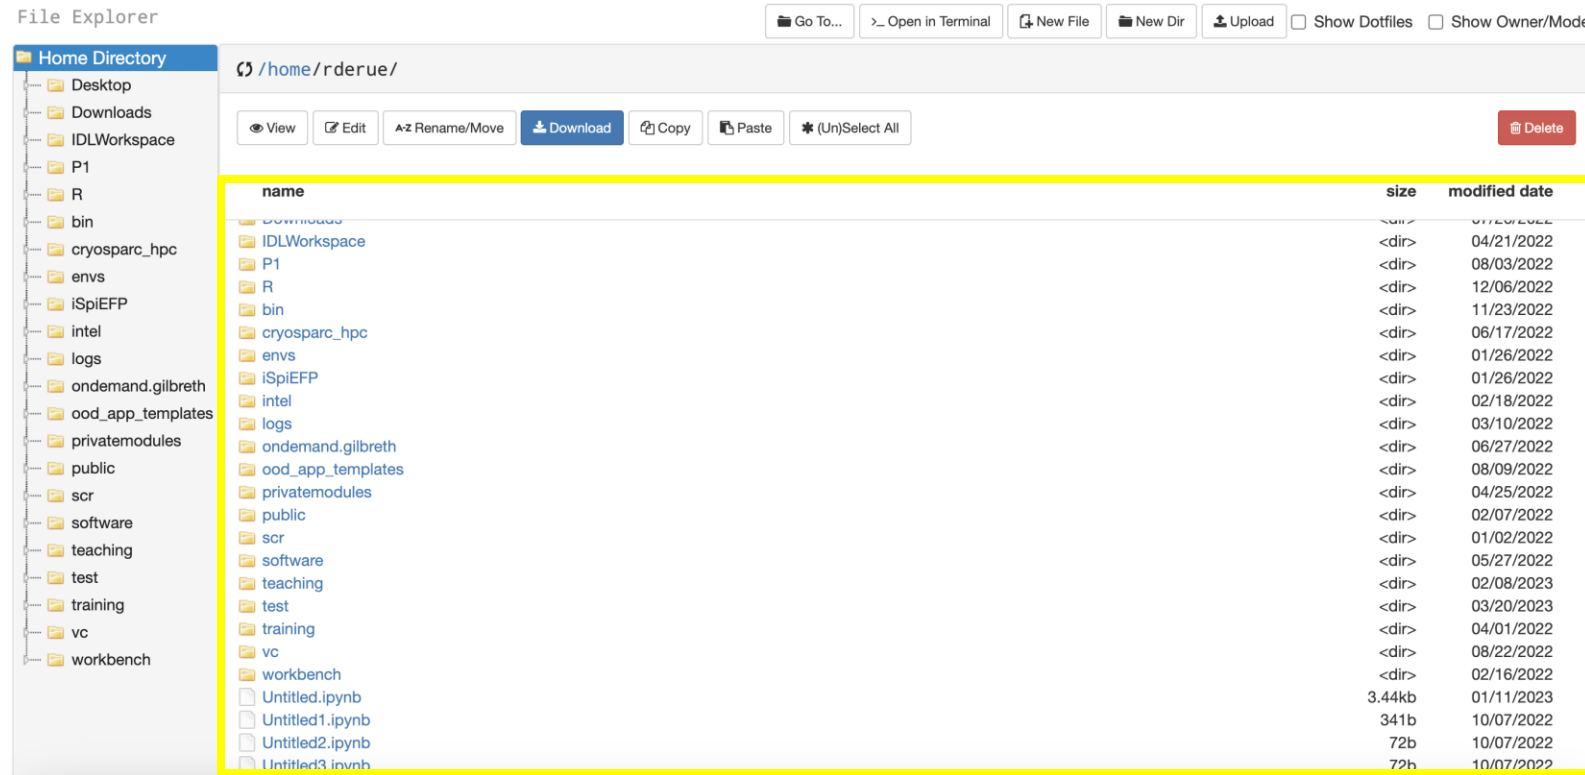
2. The menu along the left-hand side of the screen is for navigating to the different subdirectories contained within your home directory
 - Even if you change the directory you are currently examining, this menu will continue to display subdirectories relative to your home directory



Using Open OnDemand: Files

A Web-Based File Explorer

- The central part of this webpage displays the contents of your present working directory as well along with some metadata like the last modified timestamps



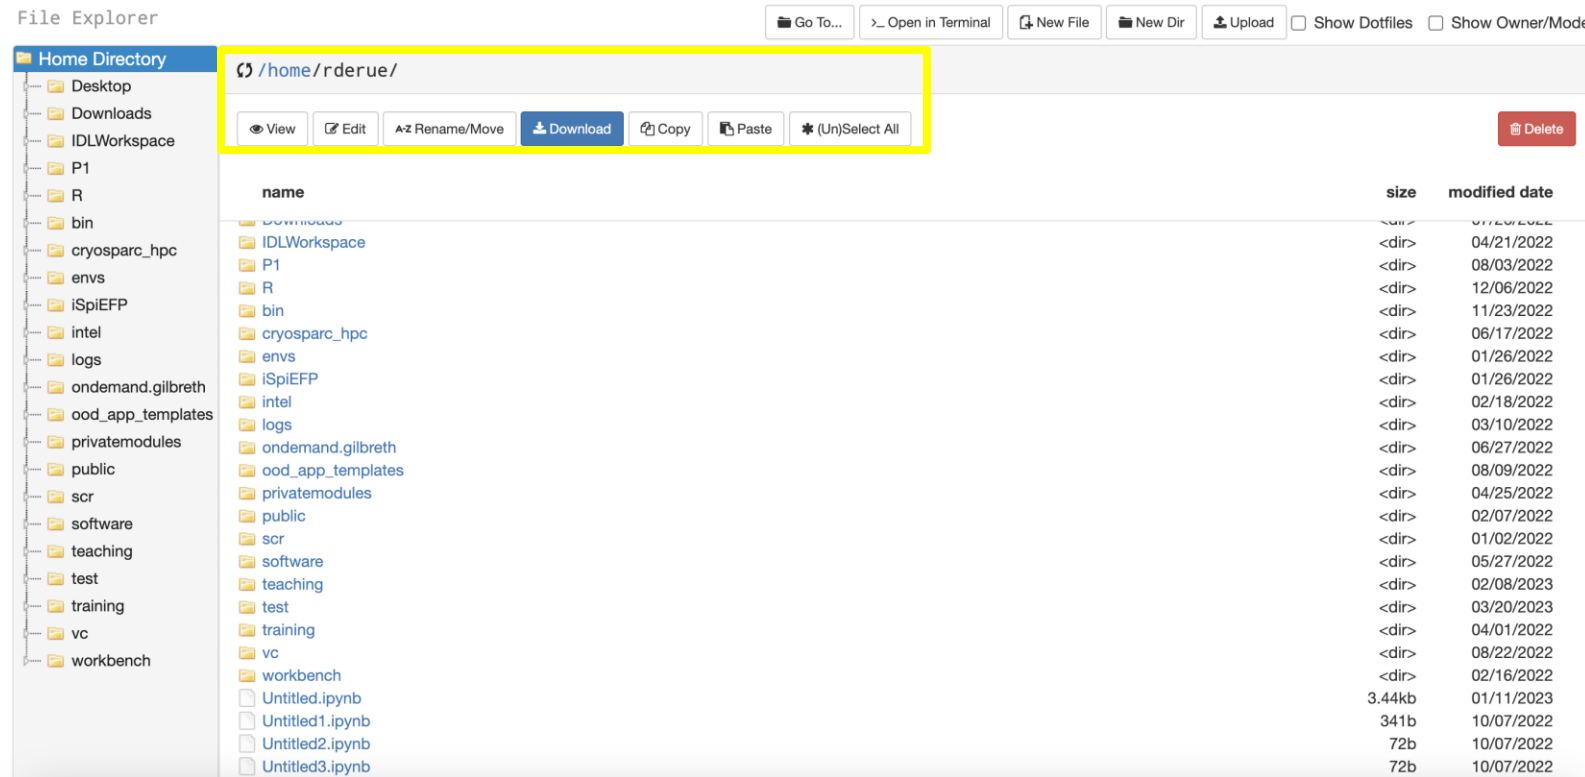
The screenshot shows a web-based file explorer interface. The left sidebar displays a tree view of the file system under 'Home Directory', including folders like Desktop, Downloads, IDLWorkspace, P1, R, bin, cryosparc_hpc, envs, iSpiEFP, intel, logs, ondemand.gilbreth, ood_app_templates, privatemodules, public, scr, software, teaching, test, training, vc, and workbench. The main area shows the contents of the current directory, which is '/home/rderue/'. The directory listing includes folders and files with their respective sizes and last modified dates.

name	size	modified date
Downloads	<dir>	07/26/2022
IDLWorkspace	<dir>	04/21/2022
P1	<dir>	08/03/2022
R	<dir>	12/06/2022
bin	<dir>	11/23/2022
cryosparc_hpc	<dir>	06/17/2022
envs	<dir>	01/26/2022
iSpiEFP	<dir>	01/26/2022
intel	<dir>	02/18/2022
logs	<dir>	03/10/2022
ondemand.gilbreth	<dir>	06/27/2022
ood_app_templates	<dir>	08/09/2022
privatemodules	<dir>	04/25/2022
public	<dir>	02/07/2022
scr	<dir>	01/02/2022
software	<dir>	05/27/2022
teaching	<dir>	02/08/2022
test	<dir>	03/20/2023
training	<dir>	04/01/2022
vc	<dir>	08/22/2022
workbench	<dir>	02/16/2022
Untitled.ipynb	3.44kb	01/11/2023
Untitled1.ipynb	341b	10/07/2022
Untitled2.ipynb	72b	10/07/2022
Untitled3.ipynb	72b	10/07/2022

Using Open OnDemand: Files

A Web-Based File Explorer

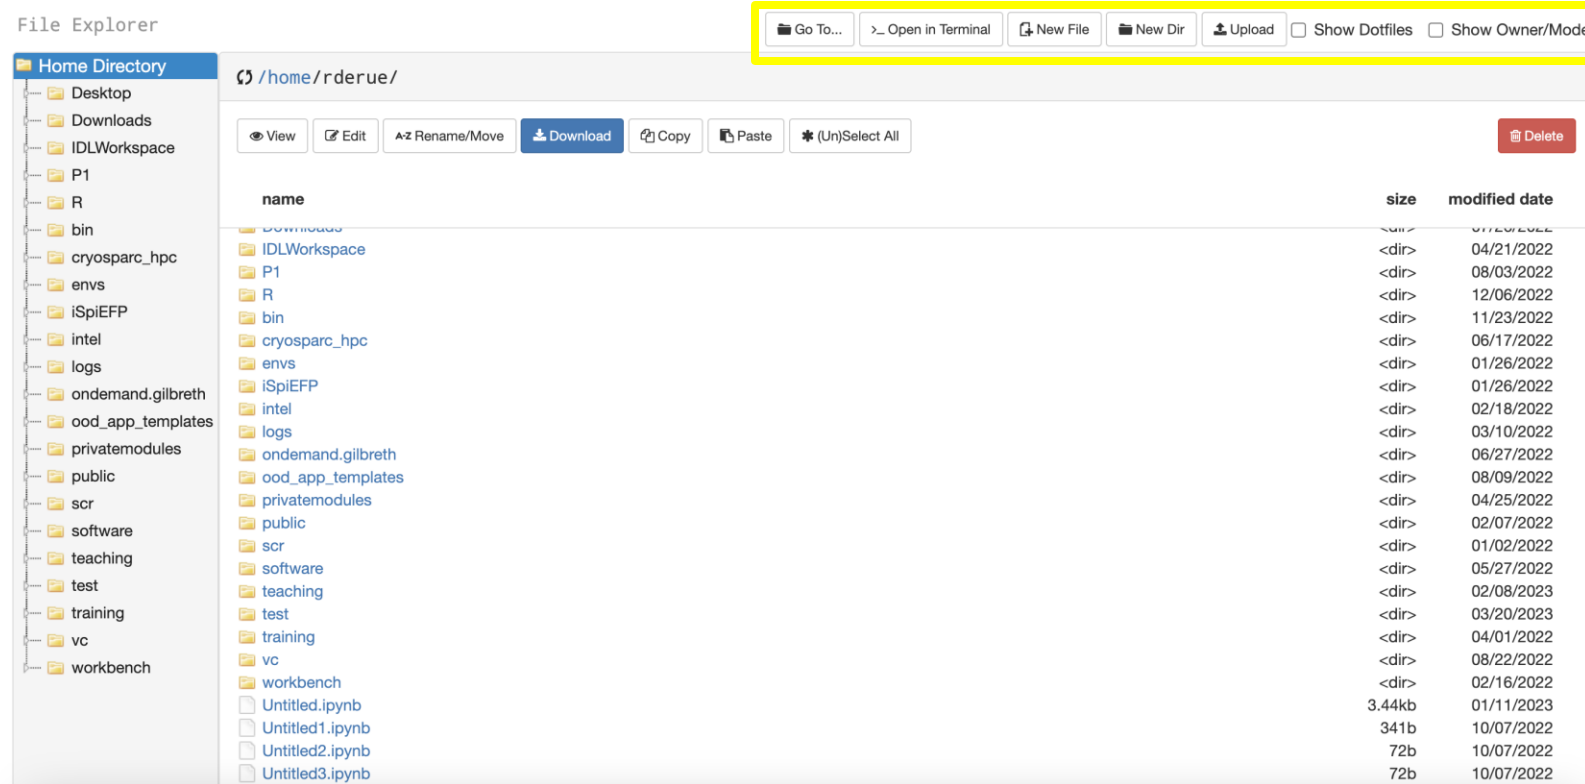
4. The toolbar located above the contents of the directory displays your current working directory as well as tools for affecting the files contained here



Using Open OnDemand: Files

A Web-Based File Explorer

- The toolbar in the upper right corner provides tools for changing your current directory, opening your current directory in a terminal to apply shell commands to the files, and creating new files/directories that are either empty or uploaded from your local workstation.



000 101

Using Open OnDemand: Applications

Using Open OnDemand: Applications

Application Forms

1. To access the App Form for Gilbreth Compute Desktop, select it from the Interactive Apps dropdown on the Dashboard
 - The form will be opened in your current tab
 - You can return to the Dashboard by clicking on "Gilbreth-Gateway" in the top left corner

The screenshot shows the Open OnDemand dashboard interface. At the top, there is a navigation bar with the following items: "Gilbreth - Gateway", "Files", "Jobs", "Interactive Apps", "Documentation", "Cluster", "My Interactive Sessions", "Develop", "Help", "Logged in as rderue", and "Log Out". The "Interactive Apps" dropdown menu is open, showing a list of applications: "Biocontainers", "CryoSPARC", "Desktops", "Gilbreth Compute Desktop", "Windows", "GUIs", "Jupyter", "MATLAB", and "Rstudio Server". The "Gilbreth Compute Desktop" option is highlighted with a yellow box. Below the navigation bar, the dashboard displays "Current Quota Usage" for three paths: "/home/rderue" (Using 2.94 GB of quota, 11%), "/scratch/gilbreth/rderue" (Using 0 Bytes of quota, 1 KB), and "/depot/lslipche" (Using 995 GB of quota, 2 TB, 48%). Each path has a progress bar and a "No file count limit." label. A note at the bottom states: "* Please allow up to 15 minutes for these numbers to update."

Using Open OnDemand: Applications

Application Forms

2. You can navigate to other Interactive Apps by selecting them from the left-hand menu

The screenshot shows the Open OnDemand interface for the 'Gilbreth Compute Desktop' application. The left-hand menu is highlighted with a yellow box, showing categories like 'Interactive Apps', 'Biocontainers', 'Desktops', 'GUIs', and 'Rstudio Server'. The main content area displays the application details, including a description, a 'Queue' dropdown set to 'standby (Max 4.0 hours)', a 'Number of hours' input field set to '0.25', a 'Number of GPUs' input field set to '1', and a 'Type of GPU' dropdown set to 'First Available'. A 'Launch' button is visible at the bottom, along with a note: '* All Gilbreth Compute Desktop session data is generated and stored under the user's'.

Using Open OnDemand: Applications

Application Forms

3. In order to request access to dedicated resources, you must select the queue to which you wish to submit your job
 - Different queues serve different purposes and information about each can be found in our user guide

Home / My Interactive Sessions / Gilbreth Compute Desktop

Interactive Apps

Biocontainers

- CryoSPARC

Desktops

- Gilbreth Compute Desktop**
- Windows

GUIs

- Jupyter
- MATLAB
- Rstudio Server

Bioinformatics Apps [Sandbox]

Biocontainers

- CryoSPARC

Interactive Apps [Sandbox]

Desktops

- Desktop

Gilbreth Compute 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.

Queue

standby (Max 4.0 hours) ▾

Please select a queue from the drop-down and enter the number of hours below (up to the max listed above).

Number of hours

0.25

Number of GPUs

1

Type of GPU

First Available ▾

I would like to receive an email when the session starts

Launch

* All Gilbreth Compute Desktop session data is generated and stored under the user's

Using Open OnDemand: Applications

Application Forms

4. You must also specify the number/type of resources you wish to have allocated and for how long you intend to use them
 - You will be given memory proportional to the number of CPUs/GPUs requested. For GPUs, typically 1 GPU is sufficient
 - It's important to remember that the scheduler prioritizes more modest requests. You should not ask for the maximum duration walltime just because it's possible

The screenshot shows the Open OnDemand interface for the 'Gilbreth Compute Desktop' application. The breadcrumb trail is 'Home / My Interactive Sessions / Gilbreth Compute Desktop'. On the left, there are three category lists: 'Interactive Apps' (CryoSPARC, Desktops: Gilbreth Compute Desktop, Windows; GUIs: Jupyter, MATLAB, Rstudio Server), 'Bioinformatics Apps [Sandbox]' (CryoSPARC), and 'Interactive Apps [Sandbox]' (Desktops: Desktop). The main content area is titled 'Gilbreth Compute Desktop' and contains the following text: '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.' Below this is a 'Queue' dropdown menu set to 'standby (Max 4.0 hours)'. A note says: 'Please select a queue from the drop-down and enter the number of hours below (up to the max listed above)'. A yellow box highlights the resource selection fields: 'Number of hours' (input: 0.25), 'Number of GPUs' (input: 1), and 'Type of GPU' (dropdown: First Available). There is a checkbox for 'I would like to receive an email when the session starts' and a blue 'Launch' button. A footer note states: '* All Gilbreth Compute Desktop session data is generated and stored under the user's'.

Using Open OnDemand: Applications

Application Forms

4. Because there can be a wait time dependent on the queue selected as well as the size/duration of the allocation requested, you can check this box to receive an e-mail when your session is ready.

Home / My Interactive Sessions / Gilbreth Compute Desktop

Interactive Apps

Biocontainers

- CryoSPARC

Desktops

- Gilbreth Compute Desktop**
- Windows

GUIs

- Jupyter
- MATLAB
- Rstudio Server

Bioinformatics Apps [Sandbox]

Biocontainers

- CryoSPARC

Interactive Apps [Sandbox]

Desktops

- Desktop

Gilbreth Compute 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.

Queue

standby (Max 4.0 hours) ▾

Please select a queue from the drop-down and enter the number of hours below (up to the max listed above).

Number of hours

0.25

Number of GPUs

1

Type of GPU

First Available ▾

I would like to receive an email when the session starts

Launch

* All Gilbreth Compute Desktop session data is generated and stored under the user's

Using Open OnDemand: Applications

Application Forms

5. After selecting the options you would like for your session, click launch to send your request to the scheduler

The screenshot shows the Open OnDemand interface for launching an application. The breadcrumb trail is 'Home / My Interactive Sessions / Gilbreth Compute Desktop'. On the left, there are three main categories: 'Interactive Apps', 'Bioinformatics Apps [Sandbox]', and 'Interactive Apps [Sandbox]'. Under 'Interactive Apps', 'Desktops' is selected, and 'Gilbreth Compute Desktop' is highlighted. Under 'Bioinformatics Apps [Sandbox]', 'CryoSPARC' is selected. Under 'Interactive Apps [Sandbox]', 'Desktops' is selected, and 'Desktop' is highlighted. The main content area is titled 'Gilbreth Compute Desktop' and contains the following text: '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.' Below this, there are several form fields: 'Queue' is a dropdown menu set to 'standby (Max 4.0 hours)'; 'Number of hours' is a text input field containing '0.25'; 'Number of GPUs' is a text input field containing '1'; 'Type of GPU' is a dropdown menu set to 'First Available'; and a checkbox labeled 'I would like to receive an email when the session starts' which is currently unchecked. At the bottom of the form is a blue 'Launch' button, which is highlighted with a yellow border. A footer note states: '* All Gilbreth Compute Desktop session data is generated and stored under the user's'.

Using Open OnDemand: Applications

Application Forms

- This will create a "card" for this session which will assign it two IDs
 - The first ID is the Slurm scheduler's job ID and is parenthesized next to the application name
 - The second ID is the Session ID and is the name of the directory where Open OnDemand will track information about the session
 - Clicking this session ID will open this directory in the File Explorer
 - When the scheduler finds resources for your job, the card will transform into the second image, and you can access the application by clicking the blue button labelled "Launch..."

Interactive Apps

Biocontainers

- CryoSPARC

Desktops

- Gilbreth Compute Desktop

Windows

GUIs

- Jupyter

Interactive Apps

Biocontainers

- CryoSPARC

Desktops

- Gilbreth Compute Desktop

Windows

GUIs

- Jupyter

Gilbreth Compute Desktop (2677465) Queued

Created at: 2023-03-24 10:02:28 EDT Delete

Time Requested: 15 minutes

Session ID: 64d864b0-246a-4b33-869a-9094c4200098

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

Gilbreth Compute Desktop (2677465) 1 node | 8 cores | Running

Host: gilbreth-f003.rcac.purdue.edu Delete

Created at: 2023-03-24 10:02:28 EDT

Time Remaining: 6 minutes

Session ID: 64d864b0-246a-4b33-869a-9094c4200098

[Launch noVNC in New Tab](#) View Only (Share-able Link)

Using Open OnDemand: Applications

Specific Applications: Compute Desktop


Every Purdue RCAC cluster has an application named "Compute Desktop."

- This application will launch a lightweight graphical desktop that you can use to interact with the HPC cluster as you would your own local workstation
- It includes a native file explorer, Firefox as an internet browser, a terminal, and a typical Desktop that you can save files to
- This application can be used to launch software that requires a GUI
 - This can be used on system-wide installations by RCAC as well as with software you personally install

Using Open OnDemand: Applications

Specific Applications: Compute Desktop


1. When you first launch the Compute Desktop, you will be presented with a desktop that appears as the one on the right
 - You can access your files by selecting the folder icon in the taskbar along the bottom



Using Open OnDemand: Applications

Specific Applications: Compute Desktop

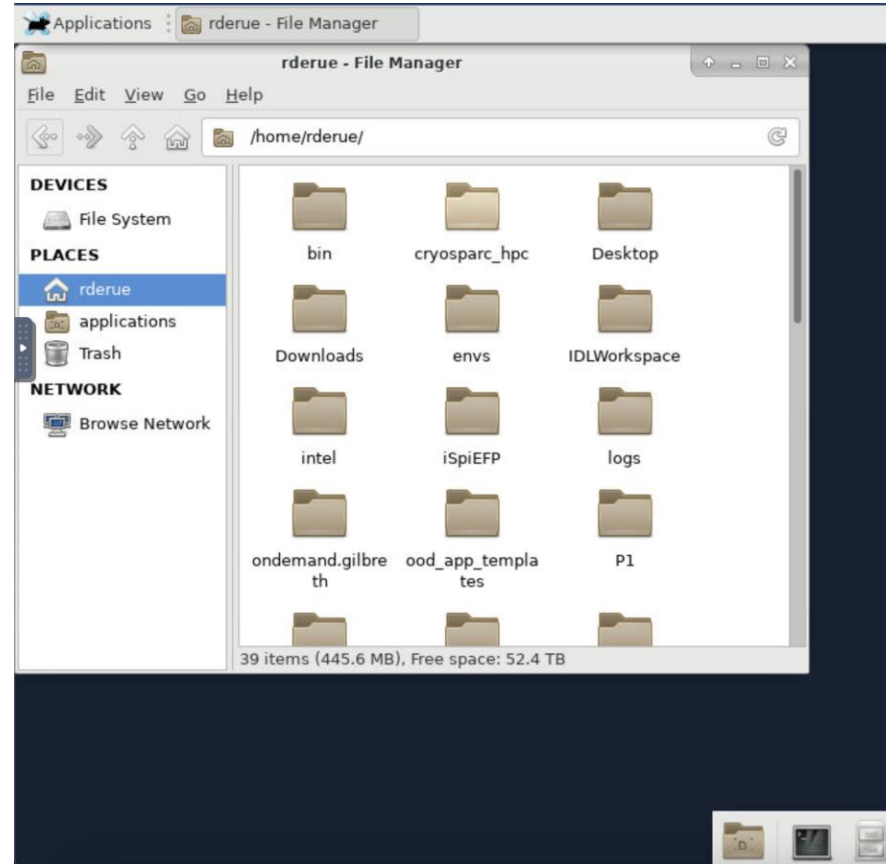
1. When you first launch the Compute Desktop, you will be presented with a desktop that appears as the one on the right
 - You can access your files by selecting the file-cabinet icon in the taskbar along the bottom



Using Open OnDemand: Applications

Specific Applications: Compute Desktop


- From this file menu, you can access any file on the cluster, but by default it will open to your home directory
 - You can change the directory you are viewing by clicking on other directories or by typing the absolute path into the search bar at the top



Using Open OnDemand: Applications

Specific Applications: Compute Desktop


2. You can launch a terminal on the desktop by selecting the terminal icon from the bottom taskbar
 - This is extremely useful for launching GUI applications from the command-line



Using Open OnDemand: Applications

Specific Applications: Compute Desktop

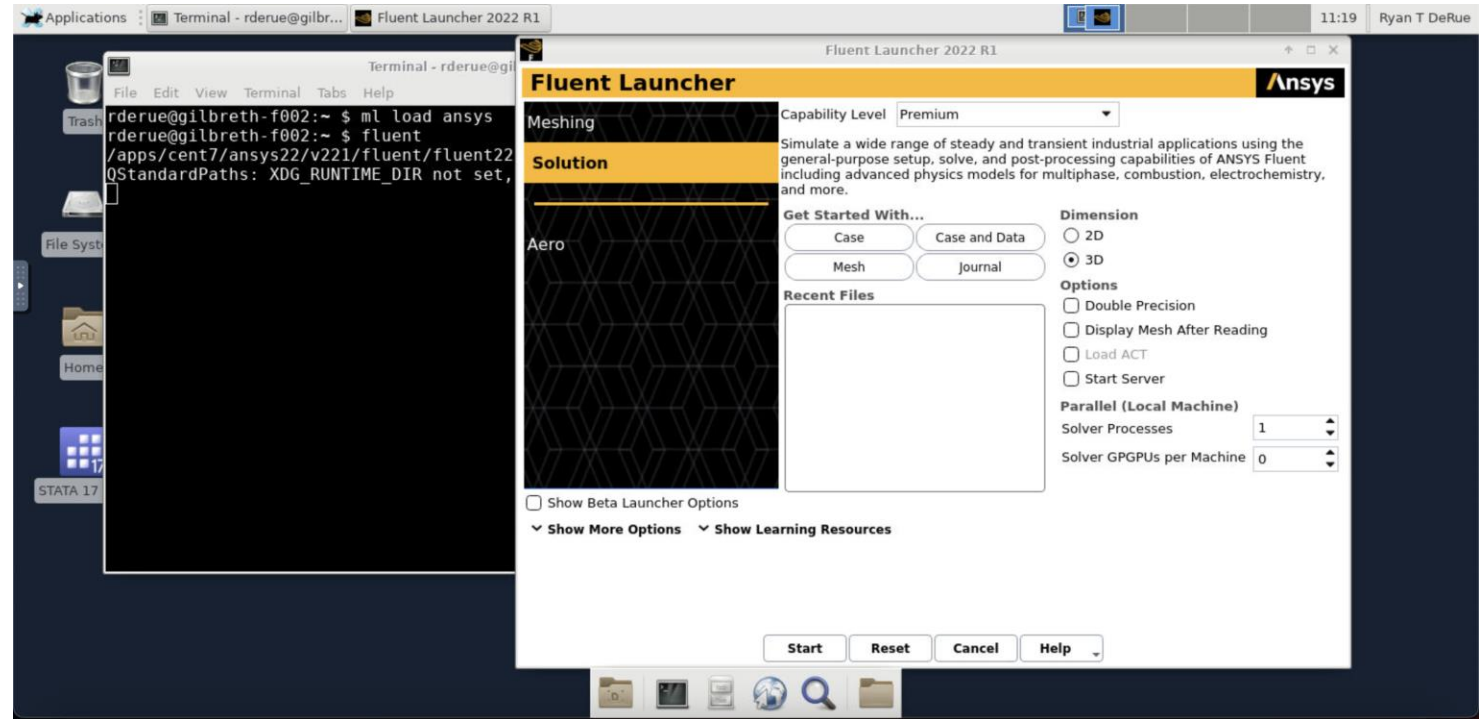
3. Notice that the prompt says that host is gilbreth-f002. This is the name of the compute node that we are on
 - If we were to try to do the same thing through ThinLinc, which provides a similar interface, the terminal would indicate that we are on a front-end (not a compute node) with a label of gilbreth-feXX



Using Open OnDemand: Applications

Specific Applications: Compute Desktop

- By loading a module (ansys) for software that I know provides a GUI, I can launch that GUI by typing the name of the command to do so.
 - Usually, you can find such commands in the user manual for software you personally install
 - This allows me to run scientific software the same way many users do on ThinLinc, but I get all the resources dedicated to my application which improves performance!



Using Open OnDemand: Applications

Specific Applications: Jupyter

1. Jupyter is by far the most popular application for Open OnDemand on our clusters
 - Very useful application for interactive work with languages like Python, R, Julia, etc.
 - App form is very similar to the previous one, but it does have an option to select whether you want to run "Jupyter Lab" or "Jupyter Notebook".
 - Notebook is the original Jupyter application and is a more bare-bones method of using Jupyter

Jupyter

This app will launch a Jupyter Notebook or Jupyter Lab session on the [Gilbreth cluster](#).

Queue

standby (Max 4.0 hours) ▾

Please select a queue from the drop-down and enter the number of hours below (up to the max listed above).

Number of hours

1

Number of GPUs

1

Type of GPU

First Available ▾

I would like to receive an email when the session starts

Launch Jupyter Lab or Jupyter Notebook

Jupyter Lab ▾

In order to use your existing Anaconda environments as Jupyter kernels within Open OnDemand, they must contain IPyKernel and IPython

Launch

Using Open OnDemand: Applications

Specific Applications: Jupyter

2. We have installed several system wide Jupyter kernels for various languages.
 - Some of these kernels, the “learning kernels” contain ready to use installations of the most common machine learning libraries
 - You cannot install additional packages into these kernels, but you can create your own anaconda environment containing the packages you want
 - If your anaconda environment contains the packages “IPyKernel” and “IPython”, you will be able to select them as a Jupyter kernel when using Open OnDemand

Jupyter

This app will launch a Jupyter Notebook or Jupyter Lab session on the [Gilbreth cluster](#).

Queue

standby (Max 4.0 hours) ▾

Please select a queue from the drop-down and enter the number of hours below (up to the max listed above).

Number of hours

1

Number of GPUs

1

Type of GPU

First Available ▾

I would like to receive an email when the session starts

Launch Jupyter Lab or Jupyter Notebook

Jupyter Lab ▾

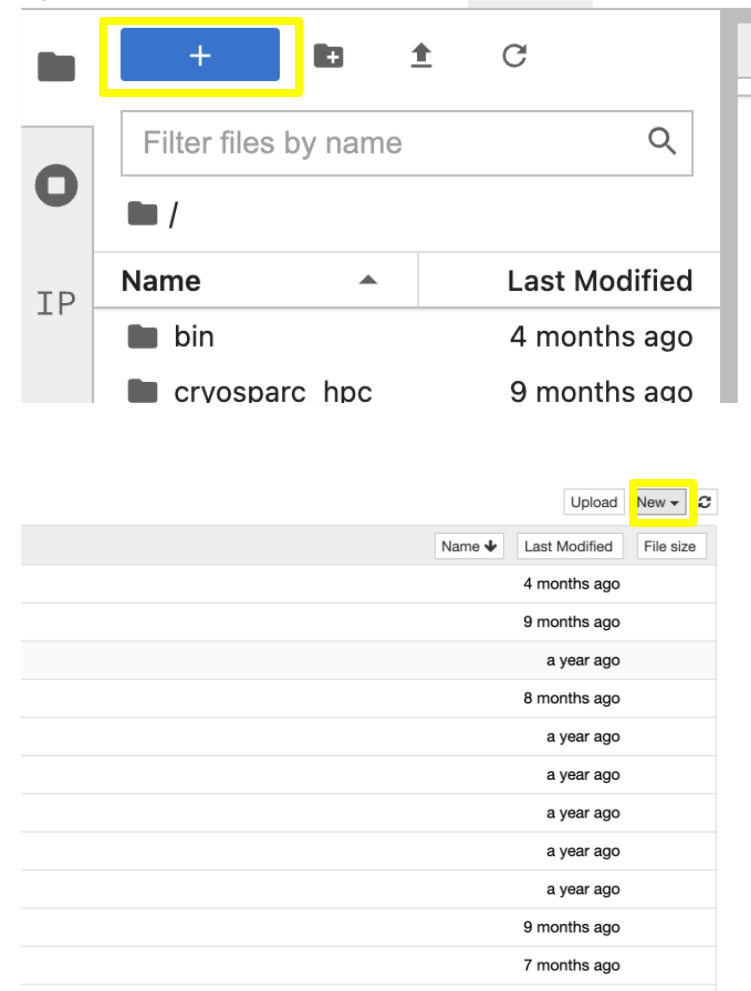
In order to use your existing Anaconda environments as Jupyter kernels within Open OnDemand, they must contain IPyKernel and IPython

Launch

Using Open OnDemand: Applications

Specific Applications: Jupyter

3. Regardless of whether you choose to use Jupyter Notebook or Jupyter Lab, it will open to your home directory.
 - To create a new notebook in Jupyter Lab click the blue “+” button in the top left corner (Shown in the top-most picture)
 - To create a new notebook in Jupyter Notebook, select the “New” button in the top right corner. (Shown in the bottom-most picture)



Using Open OnDemand: Applications

Specific Applications: Jupyter

4. After we have created a new notebook, we will be presented with our new notebook in an interface shown to the right
 - Jupyter Lab is shown in top-most picture and Jupyter Notebook shown below that
 - To select the kernel you wish to use, click on the name of the kernel shown to the right in a yellow box, and an option to select a new kernel will appear

The top screenshot shows the Jupyter Lab interface. On the left is a file browser with a search bar and a table of files. The main area is a code editor for 'Untitled.ipynb'. The kernel name 'Python 3.8 - Learning [learning/conda-2020.11-py38-gpu]' is highlighted in a yellow box. The code in the editor is:

```
[1]: import torch
[2]: torch.cuda.is_available()
[2]: True
[3]: torch.cuda.device_count()
[3]: 1
```

The bottom screenshot shows the Jupyter Notebook interface. The kernel name 'Python 3.8 - Learning [learning/conda-2020.11-py38-gpu]' is also highlighted in a yellow box. The code and output are:

```
In [1]: import torch
In [2]: torch.cuda.is_available()
Out[2]: True
In [3]: torch.cuda.device_count()
Out[3]: 1
```

000 101

Using Open OnDemand: Job Composer

Using Open OnDemand: Job Composer

A Widget For Talking to the Scheduler

Often, we may also want to run non-interactive jobs. In such cases Open OnDemand provides an interface for creating such jobs as well as monitoring your interactive and non-interactive jobs.

We can access these tools using the Jobs menu

The screenshot shows the Open OnDemand web interface. The top navigation bar includes 'Gilbreth - Gateway', 'Files', 'Jobs', 'Interactive Apps', 'Documentation', 'Cluster', 'My Interactive Sessions', 'Develop', 'Help', 'Logged in as rderue', and 'Log Out'. The 'Jobs' menu is highlighted with a yellow box, showing 'Active Jobs' and 'Job Composer' options. Below the navigation bar, the main content area displays 'OnDemand' and 'Current Quota Usage' for three paths: /home/rderue, /scratch/gilbreth/rderue, and /depot/lslipche. Each path shows the amount of quota used as a percentage of the total quota, along with a 'No file count limit' indicator and an update timestamp.

Path	Quota Used	Quota Limit	File Count Limit	Last Updated
/home/rderue	Using 2.94 GB of quota 25 GB (11%)	25 GB	No file count limit.	Updated 7 months, 1 week, 6 days, and 17 minutes ago
/scratch/gilbreth/rderue	Using 0 Bytes of quota 1 KB	1 KB	No file count limit.	No data yet
/depot/lslipche	Using 995 GB of quota 2 TB (48%)	2 TB	No file count limit.	Updated less than 1 second ago

* Please allow up to 15 minutes for these numbers to update.

Using Open OnDemand: Job Composer

A Widget For Talking to the Scheduler

1. The Active Jobs widget will allow you to query certain usage statistics from the scheduler that you would otherwise need to do through the command line interface (CLI)
 - When a job fails due to an error, this is a useful tool to discover what may have happened
 - You can look at the amount of memory that was requested to see if perhaps there was an out of memory error
 - You can also use the “Open in File Manager” button to view the output.log file that may contain information about the error that occurred

Open OnDemand / Active Jobs

Your Jobs ▾ All Clusters ▾

Active Jobs

Show 50 entries Filter:

ID	Name	User	Account	Time Used	Queue	Status	Cluster
> 2677533	sbatch	rderue	standby		gilbreth-standby	Completed	Gilbreth
> 2677510	OnDemand/Notebook	rderue	standby	1670	gilbreth-standby	Running	Gilbreth

Running OnDemand/Notebook 2677509

Cluster	Gilbreth
Job Id	2677509
Job Name	OnDemand/Notebook
User	rderue
Account	standby
Partition	gilbreth-standby
State	RUNNING
Reason	None
Total Nodes	1
Node List	gilbreth-f004
Total CPUs	8
Time Limit	1:00:00
Time Used	35:14
Memory	4096M

Output Location: `/home/rderue/ondemand.gilbreth/data/sys/dashboard/batch_connect/sys/bc_jupyter/gilbreth/output/c17f4b21-c241-4596-87a6-2752760ce9ce`

[Open in File Manager](#) [Open in Terminal](#) [Delete](#)

Using Open OnDemand: Job Composer

A Widget For Talking to the Scheduler

2. The Job Composer widget is used for submitting non-interactive jobs to the scheduler. This is very useful when you have some work to accomplish as part of your workflow that doesn't need to be done interactively.
 - An example of this might be post processing the output of some work that was generated interactively
 - This tool provides a simple default template for submitting a "Hello World" job script, but allows you to save template job scripts for workflows you may regularly repeat

The screenshot shows the Open OnDemand Job Composer interface. At the top, there is a navigation bar with "Open OnDemand / Job Composer", "Jobs", and "Templates". Below this, the "Jobs" section is displayed. It includes a "+ New Job" button and a "Create Template" button. A toolbar contains "Edit Files", "Job Options", "Open Terminal", "Submit", "Stop", and "Delete" buttons. Below the toolbar, there is a "Show 25 entries" dropdown and a search box. A table lists the jobs with columns for "Created", "Name", "ID", "Cluster", and "Status". One job is listed: "March 24, 2023 12:01pm", "(default) Simple Sequential Job", "2677533", "Gilbreth", and "Completed". At the bottom, it says "Showing 1 to 1 of 1 entries" and has "Previous", "1", and "Next" navigation buttons.

Using Open OnDemand: Job Composer

A Widget For Talking to the Scheduler

3. Selecting the "+ New Job" button will allow you to begin working on creating a new job either from one of your existing templates, or from the default hello world template.
 - This will create a new entry in your table of jobs with a status of "Not Submitted"

The screenshot displays the Open OnDemand Job Composer interface. At the top, there is a button labeled "Create a new job from the default template" with a dropdown menu open. The menu options are: "From Default Template", "From Template", "From Specified Path", and "From Selected Job". To the right of the menu is a "Create Template" button. Below the menu, there are several action buttons: "Open Terminal", "Submit", "Stop", and "Delete". A search bar is located to the right of these buttons. Below the search bar is a table with the following columns: "Created", "Name", "ID", "Cluster", and "Status". The table contains two entries:

Created	Name	ID	Cluster	Status
March 24, 2023 12:40pm	(default) Simple Sequential Job		Gilbreth	Not Submitted
March 24, 2023 12:01pm	(default) Simple Sequential Job	2677533	Gilbreth	Completed

At the bottom of the table, there is a pagination control showing "Showing 1 to 2 of 2 entries" and buttons for "Previous", "1", and "Next".

Using Open OnDemand: Job Composer

A Widget For Talking to the Scheduler

3. If we select the newly created job in the table, we can then click on the “Job Options” button to bring up configurable options for that job.
- The most important fields here are the job name, which is for your own reference, and the “Account”
 - “Account” is another name for queue. You can enter the name of one of the queues which you are able to select from the queue dropdown on any of the app forms

Job Options

Name

(default) Simple Sequential Job

Cluster

Gilbreth

Specify job script

main_job.sh

Account

standby

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

Job array specification

Job arrays are optional.

Using Open OnDemand: Job Composer

A Widget For Talking to the Scheduler

4. After saving the job options, we can scroll along the boxes on the right side of the table to do two things
 - We can see some details about the job
 - We can edit the job script that this job will run
 - This is the most important part of this workflow because this is defining what we want our non-interactive job to do
 - Select "Open Editor" to change the contents of this script

Job Details

Job Name:
(default) Simple Sequential Job

Submit to:
Gilbreth

Account:
standby

Script location:
/home/rdrue/ondemand.gilbreth/data/sys/myjobs/projects/default/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"
```

[Open Editor](#) [Open Terminal](#) [Open Dir](#)

Using Open OnDemand: Job Composer

A Widget For Talking to the Scheduler

5. Within your job script, you will want to put all the commands which you want to be ran as part of the job, as well as any options to Slurm, and then hit "Save"
 - For more information about how to do this, see the user guide page for example Slurm job scripts

```
Save /home/rderue/ondemand.gilbreth/data/sys/myjobs/projects/default/2/main_job.sh
1 #!/bin/bash
2 # JOB HEADERS HERE
3 #
4
5 echo "Going to sleep for 30 seconds, but when I wake up, I'll greet you!"
6 sleep 30
7 echo "Hello from Open OnDemand!"
```

Using Open OnDemand: Job Composer

A Widget For Talking to the Scheduler

- After we have made our edits to the job script, we can close this tab and return to the job composer page. Now all that remains is to select this job and hit submit!
- After this point the status of the job updates from "Not Submitted" to "Running", and when the job terminates, this will update finally to "Completed"

Jobs

The screenshot shows the 'Jobs' page in Open OnDemand. At the top, there are buttons for '+ New Job', 'Create Template', 'Edit Files', 'Job Options', 'Open Terminal', 'Submit', 'Stop', and 'Delete'. A 'Submit Job' tooltip is visible over the 'Submit' button. Below the buttons, there is a search bar and a 'Show 25 entries' dropdown. The main content is a table with columns: Created, Name, ID, Cluster, and Status. The first row shows a job created on March 24, 2023 at 12:40pm, named '(default) Simple Sequential Job', with ID 2677531, cluster 'Gilbreth', and status 'Not Submitted'. The second row shows a job created on March 24, 2023 at 12:01pm, named '(default) Simple Sequential Job', with ID 2677533, cluster 'Gilbreth', and status 'Completed'. At the bottom, it says 'Showing 1 to 2 of 2 entries' and has 'Previous', '1', and 'Next' navigation buttons.

Created	Name	ID	Cluster	Status
March 24, 2023 12:40pm	(default) Simple Sequential Job		Gilbreth	Not Submitted
March 24, 2023 12:01pm	(default) Simple Sequential Job	2677533	Gilbreth	Completed

This screenshot is similar to the one above, but the status of the first job has updated to 'Running'. The job details are: Created: March 24, 2023 12:40pm; Name: (default) Simple Sequential Job; ID: 2677561; Cluster: Gilbreth; Status: Running. The second job remains 'Completed'.

Created	Name	ID	Cluster	Status
March 24, 2023 12:40pm	(default) Simple Sequential Job	2677561	Gilbreth	Running
March 24, 2023 12:01pm	(default) Simple Sequential Job	2677533	Gilbreth	Completed

A Widget For Talking to the Scheduler

7. After we have made our edits to the job script, we can close this tab and return to the job composer page. Now all that remains is to select this job and hit submit!
 - After this point the status of the job updates from “Not Submitted” to “Running”, and when the job terminates, this will update finally to “Completed”

Using Open OnDemand: Job Composer

A Widget For Talking to the Scheduler

8. After the job completes, we can view the job details along the right-hand side of the page
 - We see that the "Folder Contents" has been updated with a new output file: "slurm-2677561.out"
 - If we click the link to this file, we will see that our job successfully generated the expected output:

```
Save /home/rderue/ondemand.gilbreth/data/sys/myjobs/projects/default/2/slurm-2677561.out
1 Going to sleep for 30 seconds, but when I wake up, I'll greet you!
2 Hello from Open OnDemand!
3
```

Job Details

Job Name:
(default) Simple Sequential Job

Submit to:
Gilbreth

Account:
standby

Script location:
/home/rderue/ondemand.gilbreth/data/sys/myjobs/projects/default/2

Script name:
main_job.sh

Folder Contents:
[/main_job.sh](#)
[/slurm-2677561.out](#)

000 101

Conclusions

Conclusions

- Open OnDemand helps to lower the barrier of entry to HPC for domain scientists seeking to do computational work
- We can use Open OnDemand to manage our directories as well as transfer files to/from the cluster
- Open OnDemand is especially useful for running interactive style jobs that require a graphical user interface
- We can use Open OnDemand to manage our non-interactive jobs as well

000 101

Questions?

THANK YOU

Feel free to reach out to rderue@purdue.edu with questions.

Slides are posted at: <https://www.rcac.purdue.edu/training/ood101>