

---

# Tapis Documentation

**Joe Stubbs**

**Feb 20, 2023**



# INTRODUCTION

<b>1</b>	<b>What is SCINCO?</b>	<b>1</b>
<b>2</b>	<b>Capabilities</b>	<b>3</b>
<b>3</b>	<b>Components</b>	<b>5</b>
<b>4</b>	<b>Scinco Publications</b>	<b>7</b>
<b>5</b>	<b>Quick Start Guide</b>	<b>9</b>
<b>6</b>	<b>Login to JupyterHub</b>	<b>11</b>
<b>7</b>	<b>Selecting an Image</b>	<b>13</b>
<b>8</b>	<b>Starting a Notebook</b>	<b>15</b>
<b>9</b>	<b>Administration Portal</b>	<b>17</b>
9.1	Accessing the Administration Portal . . . . .	17
9.2	Using the Administration Portal . . . . .	18
<b>10</b>	<b>User Groups</b>	<b>19</b>
10.1	Editing a User Group . . . . .	20
<b>11</b>	<b>Hub Images</b>	<b>21</b>
11.1	Adding or Editing Images . . . . .	22
<b>12</b>	<b>Hub Volume Mounts</b>	<b>23</b>
12.1	Adding or Editing Images . . . . .	23
12.2	Configuring a Host Path Mount . . . . .	23
12.3	Configuring an NFS Mount . . . . .	24
<b>13</b>	<b>Hub Admin</b>	<b>27</b>
13.1	Viewing a User's Server . . . . .	28
13.2	Stopping a Server . . . . .	28



## WHAT IS SCINCO?

The **Scinco project** at TACC provides a hosted interactive computing platform that augments and complements batch computing capabilities available today within science gateways. SCINCO combines container technology and state-of-the-art open source technologies such as Kubernetes and JupyterHub to deliver:

- A hardened, production-grade platform that can support hundreds of projects running across thousands of nodes geographically distributed across dozens of academic clusters.
- Deep integration with advanced storage and computing resources available in existing science gateways.
- Sharing and publishing features allowing analyses to be packaged with data and shared privately with collaborators or disseminated to the community at large. Diverse areas of research, including Astronomy, Biology, Climate Science, Neuroscience, and various fields of Engineering will leverage the SCINCO platform to analyze big data, implement computational models, disseminate results and train the next generation of researchers.



## CAPABILITIES

Provides users with a customized Jupyter notebook server with access to TACC's storage and compute resources. Projects currently using JupyterHub at TACC are [Designsafe-CI](#) , Hobby Eberly Dark Energy Experiment ([HETDEX](#)) and DARPA [SHADE](#).





## COMPONENTS

At its core, the Scinco architecture utilizes a customized **JupyterHub** version, an open-source, cloud-based Jupyter project that allows users to access notebook servers running on remote machines. Scinco incorporates multitenancy or logically separated instances of the platform so that individual projects can configure various aspects of the notebook servers and data made available to their users based on their individual project needs. Basic components of Scinco are:

- OAuth-based Authenticator Plugin
- JupyterHub
- Customized notebook server images
- Administrative Portal



## **SCINCO PUBLICATIONS**

- Jamthe, A.; Stubbs, J.; Packard, M.; Chuah, J.; Looney, J.; Curbelo & Gilbert, C. (2021). Enriching scientific and interactive computing with project SCINCO: JupyterHub on Kubernetes. In Gateways 2021
- Stubbs, J., Looney, J., Poindexter, M., Chalhoub, E., Zynda, G., Ferlanti, E., Vaughn, M., Fonner, J., & Dahan, M. (2020). Integrating Jupyter into Research Computing Ecosystems: Challenges and Successes in Architecting JupyterHub for Collaborative Research Computing Ecosystems. In Practice and Experience in Advanced Research Computing (pp. 91–98). Association for Computing Machinery.



## QUICK START GUIDE

To access [JupyterHub](#) hosted at TACC, you will need to perform following steps:

- Request for TACC login credentials from Tacc User Portal ([TUP](#)) This request is usually completes within 1-2 business days.
- Please email [cicsupport@tacc.utexas.edu](mailto:cicsupport@tacc.utexas.edu) in case you did not hear back.
- Once you receive TACC credentials, login to TACC [JupyterHub](#)
- Depending on whether you have access to multiple notebook images, you may see a spawner screen with list of dockerhub images to choose from.
- You may select any of the images and launch your notebook server.
- Packages can be installed by the users in the notebook as needed.
- If you need more memory for performing your scientific experiments, please contact our staff at [cicsupport@tacc.utexas.edu](mailto:cicsupport@tacc.utexas.edu) and we can discuss it further with you.

Depending on your scientific use-case, evaluate if a separate instance of JupyterHub is needed for your project. You may follow the steps below:

- Email to [cicsupport@tacc.utexas.edu](mailto:cicsupport@tacc.utexas.edu) with a brief description of your use-case and request for a meeting with our staff.
- During the meeting, explain your use-case and mention the memory and CPU requirements for running experiments.
- Mention any specific packages that should be loaded in the JupyterHub for running your experiments.
- Evaluate whether your workflow can be containerized into a docker image.



## LOGIN TO JUPYTERHUB

Users can login to Scinco JupyterHub using their TACC credentials. Once login is successful, users can spawn their notebook servers.



### Tapis Login - Tenant tacc

Username:

Password:

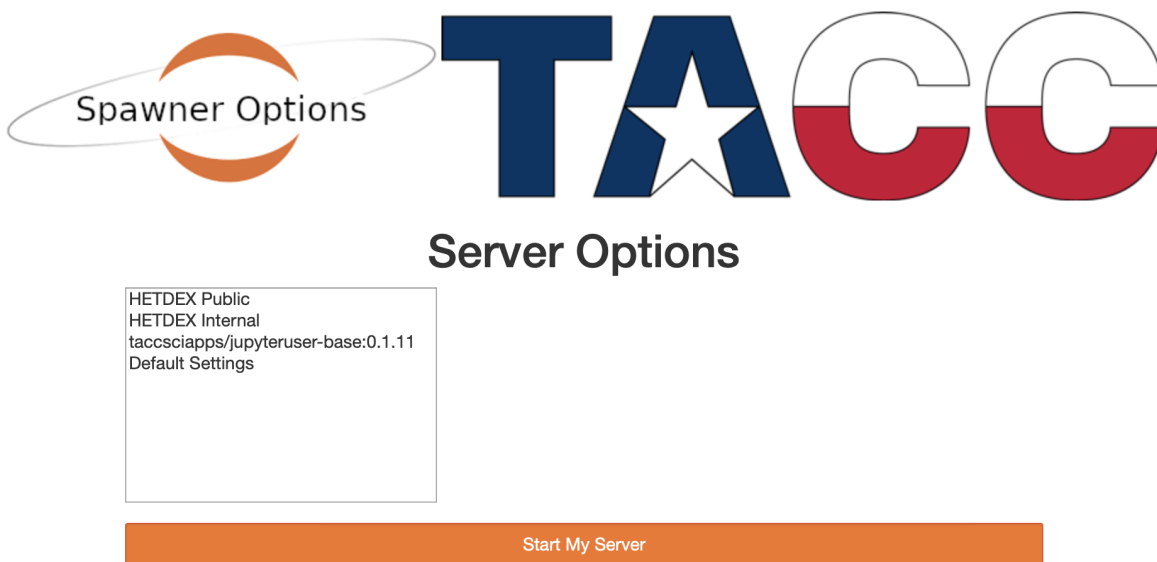
 [Forgot password?](#)





## SELECTING AN IMAGE

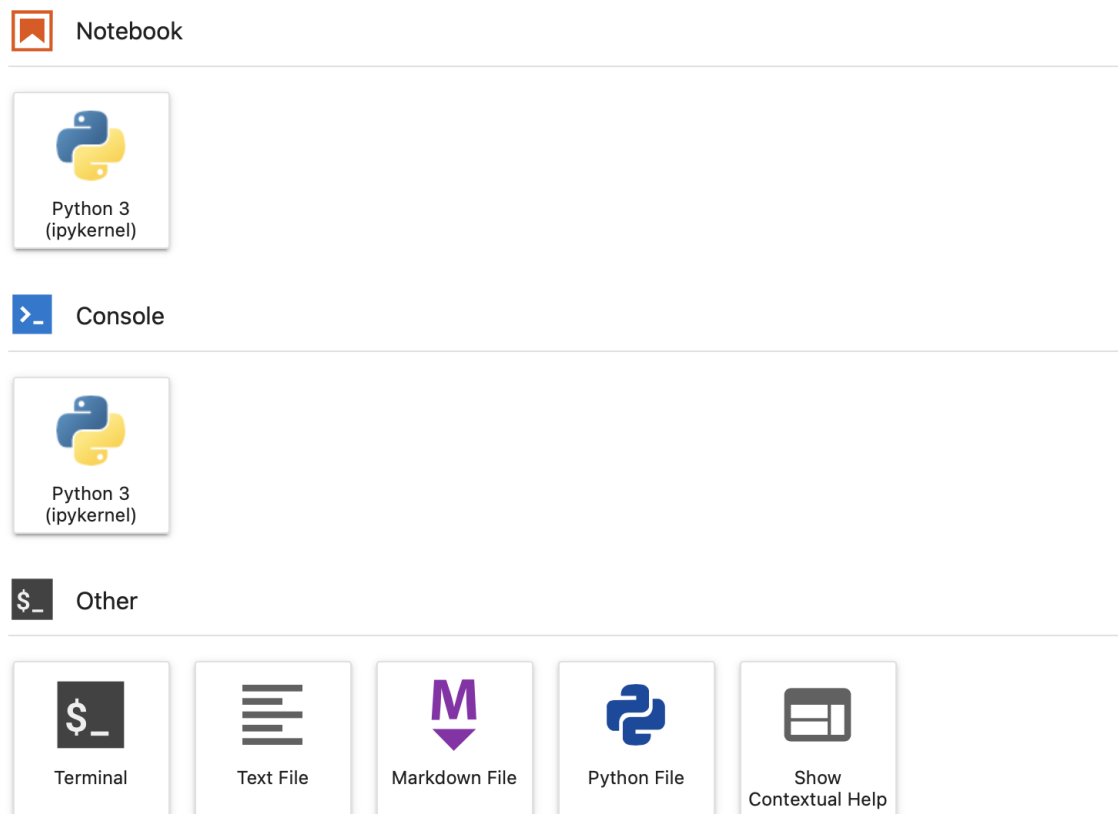
When launching a server, the JupyterHub you are using may ask you to select an image. If there is only one image, it will be selected by default. If there are multiple images available for you, you can choose one from the list.





## STARTING A NOTEBOOK

Once you have launched your server, you can run start one or more notebooks.





## **ADMINISTRATION PORTAL**

SCINCO JupyterHubs include an administrative portal. This interface allows Hub administrators to perform tasks such as:

- Starting and stopping servers on behalf of users
- Adding volume mounts to a Hub
- Adding custom images to a Hub
- See the list of users on the Hub
- Configure user groups for the Hub

If you are a PI and would like to control administrative functions on your SCINCO JupyterHub, please file a ticket for TACC staff to add you as an administrator.

### **9.1 Accessing the Administration Portal**

TACC staff can add a user to the Administration Portal instance that is running in Kubernetes. Typically, Administration Portals are available at [https://<portal\\_name>-admin.io.jupyter.tacc.cloud](https://<portal_name>-admin.io.jupyter.tacc.cloud). You will be prompted for your TACC username and password. If you have been granted access to the Administration Portal, you can log in and make alterations to the JupyterHub

## 9.2 Using the Administration Portal

Once you have successfully logged in, you will see the main page:

Home Hub ▾

# JupyterHub Admin Portal

Welcome to the JupyterHub Admin portal.

- The tenant instance is **tacc.prod**
- You can sign in to the hub at <https://jupyter.tacc.cloud>

## Hub Configuration

- [User Groups](#)
- [Hub Volume Mounts](#)
- [Hub Images](#)
- [Hub Admin](#)

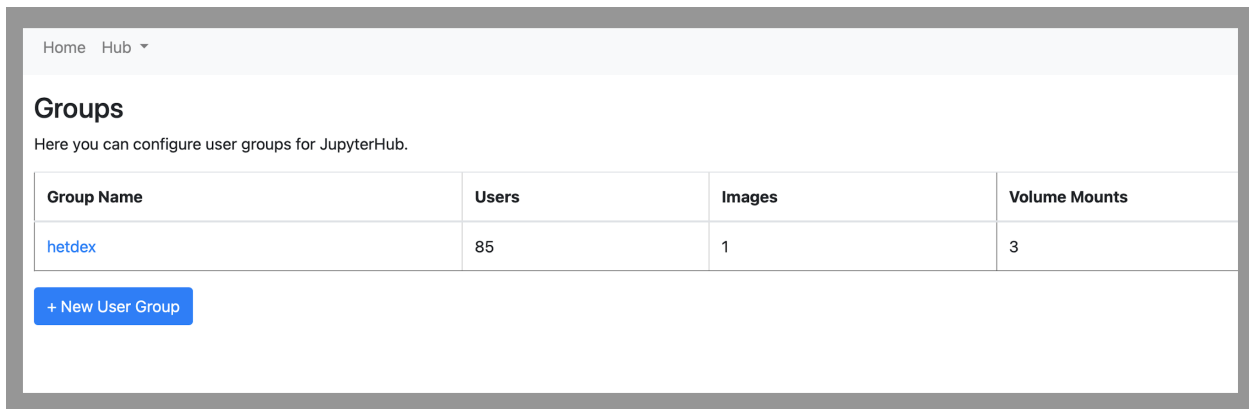
From here, you may choose the following options:

- *User Groups*
- *Hub Volume Mounts*
- *Hub Images*
- *Hub Admin*

**Warning:** The JupyterHub must be restarted by TACC staff before any configuration changes will appear.

## USER GROUPS

User Groups provide a way for JupyterHub administrators to control user access to specific server images and volume mounts based on group membership. When accessing this part of the Administration Portal, you should see a list of groups for JupyterHub.

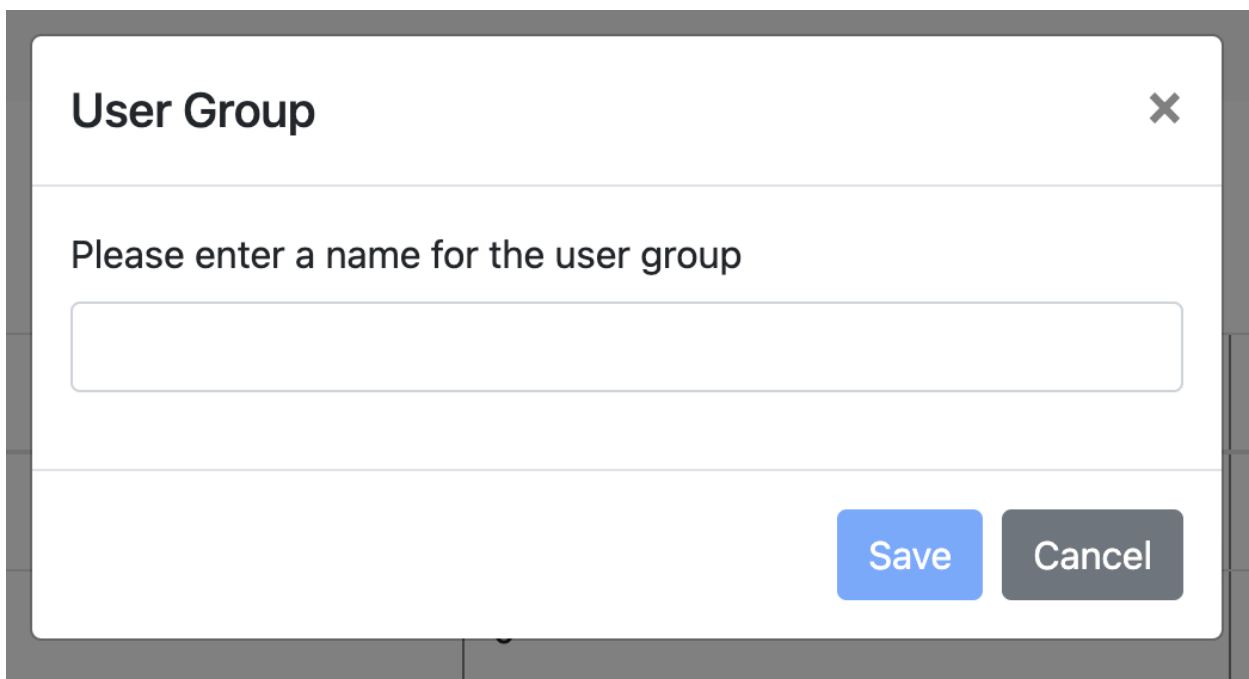


The screenshot shows the 'Groups' page in the JupyterHub Administration Portal. At the top, there is a navigation bar with 'Home' and 'Hub' (with a dropdown arrow). Below the navigation bar, the title 'Groups' is displayed, followed by the subtitle 'Here you can configure user groups for JupyterHub.' A table lists the existing groups. The table has four columns: 'Group Name', 'Users', 'Images', and 'Volume Mounts'. There is one row with the group name 'hetdex' (a blue link), 85 users, 1 image, and 3 volume mounts. Below the table is a blue button labeled '+ New User Group'.

Group Name	Users	Images	Volume Mounts
<a href="#">hetdex</a>	85	1	3

[+ New User Group](#)

You can create a new User Group by clicking the **+ New User Group** button. You will be prompted for a name for this group:



The screenshot shows a modal dialog titled 'User Group' with a close button (X) in the top right corner. The dialog contains the text 'Please enter a name for the user group' above a text input field. At the bottom right of the dialog are two buttons: 'Save' (blue) and 'Cancel' (gray).

User Group

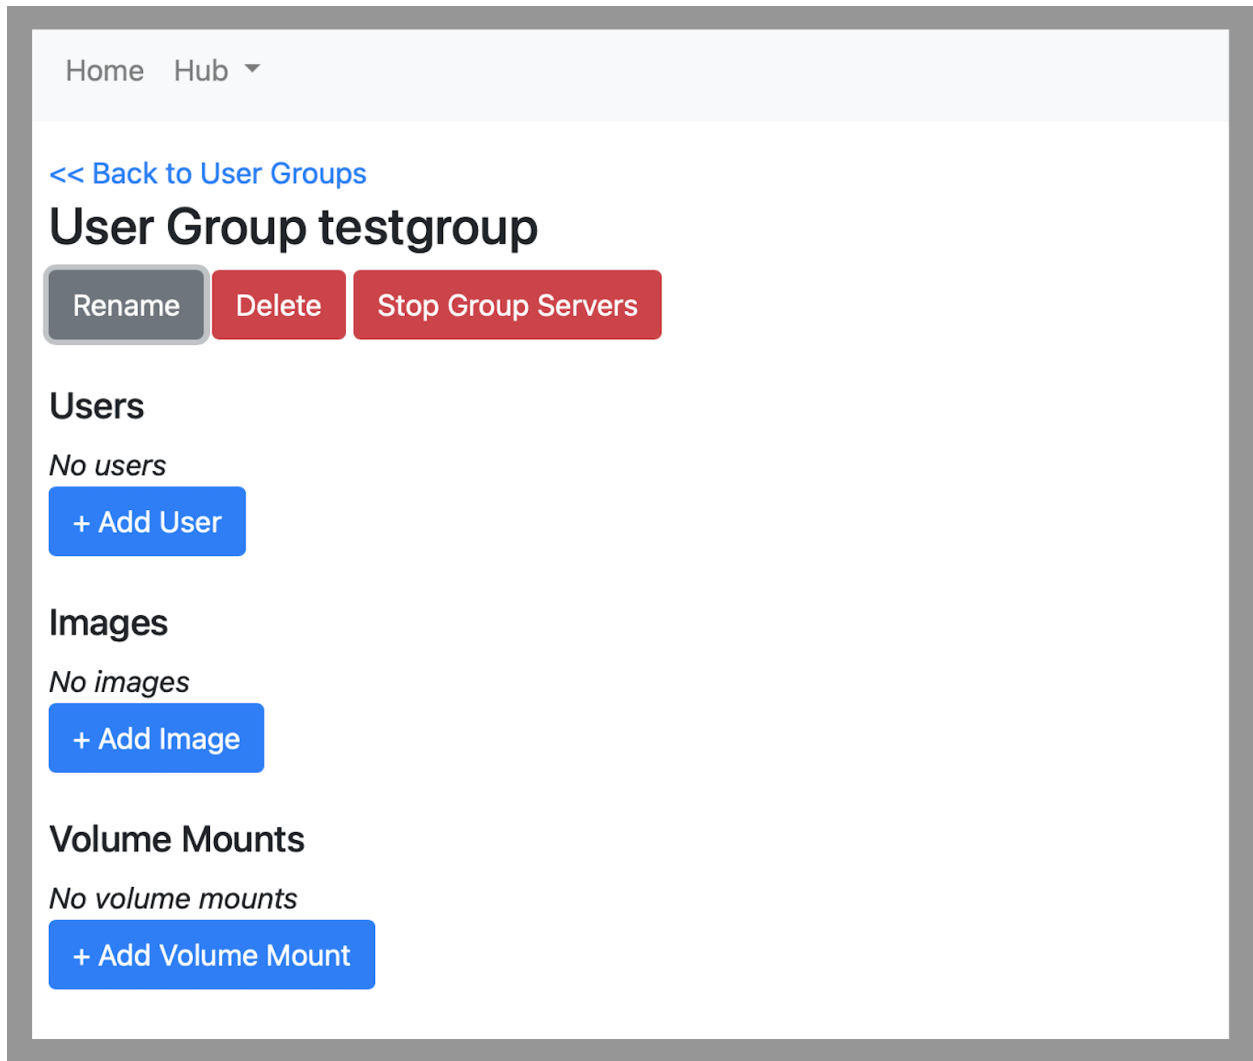
Please enter a name for the user group

Save Cancel

Alternatively, to edit an existing group you can click on the group name in the table.

## 10.1 Editing a User Group

When editing a user group, you may control who has access to this group, what images the group has in addition to global JupyterHub images, and what volume mounts will be available in addition to global JupyterHub volume mounts.



You may add a username to this group by clicking the **+ Add User** button. Alternatively, you may remove a user by clicking on the username in the group and clicking the **Delete** button.

Adding images and volume mounts follow a process that are similar to how global JupyterHub images and mounts are added. Please see the following:

- *Hub Images*
- *Hub Volume Mounts*



## HUB IMAGES

JupyterHub images are Docker images that are used to spawn a user's JupyterHub notebook server. Users may select any image that is available on your JupyterHub to spawn their notebook server. Images can be built to offer different Python packages or other customizations.

### JupyterHub Images

Here you can configure the images available to the JupyterHub spawner.

Display Name	Image Name
Default Image	taccsciapps/jupyteruser-tc:1.1.0

+ New Image

When a user logs in to a JupyterHub, **if there is more than one image available** they will be prompted to select an image to be used to start their notebook server. It may appear like this:

### Server Options

```
taccsciapps/jupyteruser-ds:1.2.24
taccsciapps/jupyteruser-sim:1.1
```

Start My Server

## 11.1 Adding or Editing Images

To add a new image option, click the + **New Image** button. Alternatively, to edit or delete an existing image option, click on its name in the table.

### JupyterHub Image Configuration

Configuration for Default Image

Display Name

Image Name

When editing an image option, the **Display Name** is the name that is shown in the Server Options prompt in the server spawner. The **Image Name** is the Docker image name used to for this server option.

## HUB VOLUME MOUNTS

Volume mounts allow a user to access file system data from within their JupyterHub server. Volume mounts can either be data that is locally available on the Kubernetes node running the server or an external NFS mount.

### JupyterHub Mounts

Here you can configure the volume mounts for the JupyterHub

Mount Path	Host Path
<a href="#">/user/mount/path</a>	/localpath/mydir

+ New Image

### 12.1 Adding or Editing Images

To add a new mount, click the **+ New Mount** button. Alternatively, to edit or delete an existing mount, click on its name in the table. When creating a mount, select a **Mount Type** of either **Host Path** or **NFS**.

### 12.2 Configuring a Host Path Mount

A Host Path mount is a Notebook Server mount to a file system location on the node that is hosting the server. For host paths to be available as mount options, they must be configured by TACC staff. If you would like a new host path mount, you must file a ticket requesting the host path before making configuration changes in the Administration Portal.

# JupyterHub Mount Configuration

Configuration for /user/mount/path

Mount Type Host Path ▾

Remote Server

The hostname of the remote server for this mount

Path

/localpath/mydir

Mount Path

/user/mount/path

☒ Read Only

Save

Delete

The **Remote Server** option is disabled for Host Path mounts. The **Path** option is the file system path on the host node that will be mounted. The **Mount Path** option is the path as it appears in a user's Notebook server. The **Read Only** option restricts a user's Notebook server write access to that path.

## 12.3 Configuring an NFS Mount

An NFS mount is a Notebook Server mount to a remote file system that has been exported via NFS. If you would like to configure an NFS mount, you must file a ticket to verify that an NFS export has been configured before making configuration changes in the Administration Portal.

# JupyterHub Mount Configuration

Configuration for /user/mount/path

Mount Type NFS ▼

Remote Server

remote.server

Path

/localpath/mydir

Mount Path

/user/mount/path

☒ Read Only

Save

Delete

The configuration options here are similar to those available when configuring a Host Path mount, except that you may now specify an address for the **Remote Server**. It does not need to be prefixed with `nfs://`.



## HUB ADMIN

When a user logs in to a JupyterHub the first time, they are registered as a user on the JupyterHub. At this point, their activity can be viewed through **Hub Admin** and you may perform administrative actions on their Notebook server.

### JupyterHub Servers

This is a list of users on the JupyterHub. If they currently have a running server, you can view it and shut it down. Otherwise, the user is listed as *inactive*.

Username ▼	Started ▼	Last Activity ▼
abc	<i>Inactive</i>	<i>Inactive</i>
xyz	2021-06-30 19:11	2021-07-01 18:42

You may sort the list of users by clicking the arrows at the top of each column on the table. Users that are listed as *Inactive* do not have a running Notebook server. Users that have running Notebook servers will show timestamps for their activity. To view a user, click on their name. When viewing an active user, you will see a screen that may look like this:

### JupyterHub User xyz

#### Started

2021-06-30 19:11

#### Last Activity

2021-07-01 18:42

[View User Server](#)

[Stop Server](#)

## 13.1 Viewing a User's Server

By clicking the **View User Server** button, you will be taken to the user's active Notebook server. You are logged in to their server *as that user* using a token for authorization. From here, you may perform any action normally available to that user on their behalf, such as opening terminals, starting notebooks or shutting down the server.

## 13.2 Stopping a Server

Sometimes it is necessary to stop a Notebook server on behalf of the user. By clicking the **Stop Server** button, the user's Notebook server will shut down. It may take some time for the operation to complete.