



Red Hat OpenShift Data Science 1

Managing users and user resources

Learn to manage user permissions and environments in Red Hat OpenShift Data Science

Red Hat OpenShift Data Science 1 Managing users and user resources

Learn to manage user permissions and environments in Red Hat OpenShift Data Science

Legal Notice

Copyright © 2021 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

Learn to manage user permissions and environments in Red Hat OpenShift Data Science.

Table of Contents

PREFACE	3
CHAPTER 1. USER TYPES AND PERMISSIONS	4
CHAPTER 2. ADDING USERS FOR OPENSIFT DATA SCIENCE	5
2.1. ADDING EXISTING USER GROUPS FROM AN IDENTITY PROVIDER TO OPENSIFT DATA SCIENCE	5
2.2. ADDING USERS FOR OPENSIFT DATA SCIENCE USING DEFAULT USER GROUPS	6
2.3. ADDITIONAL RESOURCES	8
CHAPTER 3. VIEWING OPENSIFT DATA SCIENCE USERS	9
CHAPTER 4. DELETING USERS AND USER RESOURCES	10
4.1. BACKING UP STORAGE DATA FROM AMAZON EBS	10
4.2. STOPPING NOTEBOOK SERVERS OWNED BY OTHER USERS	11
4.3. REVOKING USER ACCESS TO JUPYTERHUB	12
4.4. CLEANING UP AFTER DELETING USERS	13
CHAPTER 5. ALLOCATING ADDITIONAL RESOURCES TO OPENSIFT DATA SCIENCE USERS	15
CHAPTER 6. MANAGING NOTEBOOK SERVERS	16
6.1. ACCESSING THE JUPYTERHUB ADMINISTRATION INTERFACE	16
6.2. STARTING NOTEBOOK SERVERS OWNED BY OTHER USERS	16
6.3. ACCESSING NOTEBOOK SERVERS OWNED BY OTHER USERS	17
6.4. STOPPING NOTEBOOK SERVERS OWNED BY OTHER USERS	17
CHAPTER 7. BACKING UP STORAGE DATA FROM AMAZON EBS	19

PREFACE

This documentation is provided for the Field Trial release of Red Hat OpenShift Data Science.

See the following documents for service and life cycle information related to this Field Trial release:

- [OpenShift Data Science Service Definition](#)
- [OpenShift Data Science Life Cycle](#)


CHAPTER 1. USER TYPES AND PERMISSIONS

Red Hat OpenShift Data Science uses different user groups to control the permissions available to each user.

Red Hat OpenShift Data Science contains the following user types:

Table 1.1. User types and permissions

User Type	Default User Group	Permissions
Data scientists	rhods-users	Data scientists can access and use individual components of Red Hat OpenShift Data Science, such as JupyterHub.
IT operations administrators	rhods-admins	In addition to the actions permitted to a data scientist, IT operations administrators can: <ul style="list-style-type: none"> • Configure Red Hat OpenShift Data Science settings. • Access and manage notebook servers in the JupyterHub administration interface.



IMPORTANT

Although users of OpenShift Data Science and its components are authenticated through OpenShift, session management is separate from authentication. This means that logging out of OpenShift Dedicated or OpenShift Data Science does not affect a logged in JupyterHub session running on those platforms. This means that when a user's permissions change, that user must log out of all current sessions in order for the changes to take effect.



IMPORTANT

The user groups configured in OpenShift Dedicated, **cluster-admins** and **dedicated-admins**, are separate to the OpenShift Data Science user groups.

There are some operations relevant to OpenShift Data Science that require the **cluster-admins** or **dedicated-admins** role. Those operations include:

- Adding users to the **rhods-users** and **rhods-admins** groups.
- Removing users from the **rhods-users** and **rhods-admins** groups.
- Managing custom environment and storage configuration for users in OpenShift Dedicated, such as Jupyter notebook resources, ConfigMaps, and persistent volume claims (PVCs).

Additional resources

- [Administering your OpenShift Dedicated cluster](#)

CHAPTER 2. ADDING USERS FOR OPENSIFT DATA SCIENCE

You can grant users permission to access Red Hat OpenShift Data Science by adding user accounts to the Red Hat OpenShift Data Science user group, administrator group, or both. You can either use the default group name, or specify a group name that already exists in your identity provider.

The **user group** provides the user with access to developer functions in the Red Hat OpenShift Data Science dashboard, and associated services, such as JupyterHub. The default user group name is **rhods-users**.

The **administrator group** provides the user with access to developer and administrator functions in the Red Hat OpenShift Data Science dashboard and associated services, such as JupyterHub. The default administrator group name is **rhods-admins**.

To use the default group names, see *Adding users for OpenShift Data Science using default user groups*. This method is easy to set up, but you must manage the user lists manually in the OpenShift Dedicated web console.

To use groups that already exist in your identity provider, see *Adding existing user groups from an identity provider to OpenShift Data Science*. With this method you can manage users through your identity provider as you normally would.



IMPORTANT

If you are using LDAP as your identity provider, you need to configure LDAP syncing to OpenShift Dedicated. See [Syncing LDAP groups](#) for more information.

2.1. ADDING EXISTING USER GROUPS FROM AN IDENTITY PROVIDER TO OPENSIFT DATA SCIENCE

You can grant a user access to Red Hat OpenShift Data Science by adding their user name to the OpenShift Data Science user group, administrator group, or both. Follow the steps in this section to use an existing group from your identity provider that does not use one of the default group names, **rhods-admins** or **rhods-users**. You can add users to these groups as you normally would with that identity provider.

Prerequisites

- You have configured a supported identity provider for OpenShift Dedicated.
- You are part of the **cluster-admins** or **dedicated-admins** user group in OpenShift Dedicated.

Procedure

1. In the OpenShift Dedicated web console, change into the **Administrator** perspective.
2. Click **Workloads** → **ConfigMaps**.
3. Set the **Project** to **All Projects** or **redhat-ods-applications** to ensure you can see the appropriate ConfigMap.
4. Click the name of the **rhods-groups-config** ConfigMap.
The **ConfigMap details** page appears.
5. Click the **YAML** tab.

6. Change the **opendatahub.io/modified** label to **'true'**.

```
labels:  
  app: jupyterhub  
  opendatahub.io/modified: 'true'
```

7. Replace default values with your group names.
Change the value of **admin_groups** to the new name of your admin group and the value of **allowed_groups** to the new name of your user group, for example:

```
data:  
  admin_groups: it-ops  
  allowed_groups: datasci-devs1,datasci-devs2
```

8. Click **Save**.
9. Apply the new application configuration.
 - a. Change into the **Developer** perspective.
 - b. Click **Topology** and click on the JupyterHub application.
 - c. Click **Actions** → **Start Rollout** to deploy JupyterHub with its updated user configuration.

Verification

- Click the **Details** tab and confirm that the **Labels** field contains **opendatahub.io/modified: 'true'**, and the updated group names appear under the **Data** heading.
- The user can access the Red Hat OpenShift Data Science dashboard, and associated services, such as JupyterHub.

Additional resources

- [Sharing the instance address with users](#)

2.2. ADDING USERS FOR OPENSIFT DATA SCIENCE USING DEFAULT USER GROUPS

You can grant a user access to Red Hat OpenShift Data Science by adding their user name to the OpenShift Data Science user group, administrator group, or both. Follow the steps in this section to create administrator and user groups that use the default group names, and manually add users to the groups. This method is easy to set up, but you must manage the user lists manually in the OpenShift Dedicated web console.

Prerequisites

- You have configured a supported identity provider for OpenShift Dedicated.
- You are part of the **dedicated-admins** user group in OpenShift Dedicated.

Procedure

1. In the OpenShift Dedicated web console, click **User Management** → **Groups**.

2. **Optional:** If not present, create the **rhods-admins** group.

- a. Click **Create Group**.
- b. Change the **name** of the group to **rhods-admins**.

```
apiVersion: user.openshift.io/v1
kind: Group
metadata:
  name: rhods-admins
users:
  - user1
  - user2
```

- c. Skip to step 6 to add administrative users.

3. **Optional:** If not present, create the **rhods-users** group.

- a. Click **Create Group**.
- b. Change the **name** of the group to **rhods-users**.

```
apiVersion: user.openshift.io/v1
kind: Group
metadata:
  name: rhods-users
users:
  - user1
  - user2
```

- c. Skip to step 6 to add normal users.

4. Click the name of the group you want to add users to.

- For administrative users, click **rhods-admins**.
- For normal users, click **rhods-users**.

The **Group details** page for that group appears.

5. Click the **YAML** tab.

6. In the **users** section, add the user name of the user that you want to add to the group. For example:

```
users:
  - jdoe
  - emustermann
```

7. Click **Save**.

Verification

- Click the **Details** tab for each group and confirm that the **Users** section contains the user names that you added.

- Ensure the user can access the Red Hat OpenShift Data Science dashboard, and associated services, such as JupyterHub.

Additional resources

- [Sharing the instance address with users](#)

2.3. ADDITIONAL RESOURCES

- [User types and permissions](#)

CHAPTER 3. VIEWING OPENSIFT DATA SCIENCE USERS

You can view users who have permission to access Red Hat OpenShift Data Science. Users permitted to access Red Hat OpenShift Data Science belong to the Red Hat OpenShift Data Science user group, administrator group, or both.

Prerequisites

- The Red Hat OpenShift Data Science user group, administrator group, or both exist.
- You are part of the **dedicated-admins** user group in OpenShift Dedicated.
- You have configured a supported identity provider for OpenShift Dedicated.

Procedure

1. In the OpenShift Dedicated web console, click **User Management** → **Groups**.
2. Click the name of the group containing the users that you want to view.
 - For administrative users, click **rhods-admins**.
 - For normal users, click **rhods-users**.

The **Group details** page for the group appears.

Verification

- In the **Users** section for the relevant group, you can view the users who have permission to access Red Hat OpenShift Data Science.

CHAPTER 4. DELETING USERS AND USER RESOURCES

Users with administrator access to OpenShift Dedicated can revoke user access to JupyterHub and delete user resources from Red Hat OpenShift Data Science.



IMPORTANT

To completely remove a user from OpenShift Data Science, you must remove them from the allowed group in your OpenShift identity provider.

4.1. BACKING UP STORAGE DATA FROM AMAZON EBS

Red Hat recommends that you back up the data on your persistent volume claims (PVCs) regularly. Backing up your data is particularly important before deleting a user and before uninstalling OpenShift Data Science, as all PVCs are deleted when OpenShift Data Science is uninstalled.

Prerequisites

- You have credentials for OpenShift Cluster Manager (<https://console.redhat.com/openshift/>).
- You have administrator access to the OpenShift Dedicated cluster.
- You have credentials for the Amazon Web Services (AWS) account that the OpenShift Dedicated cluster is deployed under.

Procedure

1. Determine the IDs of the persistent volumes (PVs) that you want to back up.
 - a. In the OpenShift Dedicated web console, change into the **Administrator** perspective.
 - b. Click **Home** → **Projects**.
 - c. Click the **rhods-notebooks** project.
The **Details** page for the project opens.
 - d. Click the **PersistentVolumeClaims** in the **Inventory** section.
The **PersistentVolumeClaims** page opens.
 - e. Note the ID of the persistent volume (PV) that you want to back up.



NOTE

The persistent volumes (PV) that you make a note of are required to identify the correct EBS volume to back up in your AWS instance.

2. Locate the EBS volume containing the PVs that you want to back up.
See [Create Amazon EBS snapshots](#) for more information.
 - a. Log in to AWS (<https://aws.amazon.com>) and ensure that you are viewing the region that your OpenShift Dedicated cluster is deployed in.
 - b. Click **Services**.

- c. Click **Compute** → **EC2**.
 - d. Click **Elastic Block Storage** → **Volumes** in the side navigation.
The **Volumes** page opens.
 - e. In the search bar, enter the ID of the persistent volume (PV) that you made a note of earlier.
The **Volumes** page reloads to display the search results.
 - f. Click on the volume shown and verify that any **kubernetes.io/created-for/pvc/namespace** tags contain the value **rhods-notebooks**, and any **kubernetes.io/created-for/pvc/name** tags match the name of the persistent volume that the EC2 volume is being used for, for example, **jupyterhub-nb-user1-pvc**.
3. Back up the EBS volume that contains your persistent volume (PV).
 - a. Right-click on the volume that you want to back up and select **Create Snapshot** from the list.
The **Create Snapshot** page opens.
 - b. Enter a **Description** for the volume.
 - c. Click **Create Snapshot**.
The snapshot of the volume is created.
 - d. Click **Close**.

Verification

- The snapshot that you created is visible on the **Snapshots** page in AWS.

Additional resources

- [Create Amazon EBS snapshots](#)

4.2. STOPPING NOTEBOOK SERVERS OWNED BY OTHER USERS

Administrators can stop notebook servers that are owned by other users to reduce resource consumption on the cluster, or as part of removing a user and their resources from the cluster.

Prerequisites

- You are part of the OpenShift Data Science administrator group (by default, **rhods-admins**) in OpenShift Dedicated.
- You have logged in to JupyterHub.
- The notebook server that you want to stop is running (started).

Procedure

1. In the JupyterHub interface, click the **Admin** tab.
2. Stop one or more servers.
 - If you want to stop one or more specific servers:

- i. Locate the user that the notebook server belongs to.
 - ii. Click the **Stop server** button beside the user.
- If you want to stop all servers:
 - i. Click the **Stop all** button.
 - ii. Click **OK** to confirm stopping all servers.

Verification

- The **Stop server** button beside each server changes to a **Start server** button when the notebook server has stopped.

4.3. REVOKING USER ACCESS TO JUPYTERHUB

You can revoke a user's access to JupyterHub to prevent them from running notebook servers and consuming resources in your cluster through JupyterHub, while still allowing them access to OpenShift Data Science and other services that use OpenShift's identity provider for authentication.



IMPORTANT

To completely remove a user from OpenShift Data Science, you must remove them from the allowed group in your OpenShift identity provider.

Prerequisites

- You have stopped any notebook servers owned by the user you want to delete.
- You are part of the **dedicated-admins** user group in OpenShift Dedicated.
- The user is part of the OpenShift Data Science user group, administrator group, or both.

Procedure

1. In the OpenShift Dedicated web console, click **User Management** → **Groups**.
2. Click the name of the group that you want to remove the user from.
 - For administrative users, click **rhods-admins**.
 - For normal users, click **rhods-users**.

The **Group details** page for the group appears.

3. In the **Users** section on the **Details** tab, locate the user that you want to remove.
4. Click the action menu (**⋮**) beside the user that you want to remove and click **Remove user**.

Verification

- Check the **Users** section on the **Details** tab and confirm that the user that you removed is not visible.
- In the **rhods-notebooks** project, check under **Workload** → **Pods** and ensure that there is no

notebook server pod for this user. If you can see a pod named **jupyterhub-nb-<username>*** for the user that you have removed, delete that pod to ensure that the deleted user is not consuming resources on the cluster.



4.4. CLEANING UP AFTER DELETING USERS

After removing a user's access to Red Hat OpenShift Data Science or JupyterHub, you must also delete their associated configuration files from OpenShift Dedicated. It is recommended that you back up the user's data and profile before removing their configuration files.

Prerequisites

- (Optional) If you want to completely remove the user's access to OpenShift Data Science, you have removed their credentials from your identity provider.
- You have revoked the user's access to JupyterHub.
- You have backed up the user's storage data from Amazon EBS.
- You are part of the **dedicated-admins** user group in OpenShift Dedicated.
- You are part of the **rhods-admins** user group in OpenShift Dedicated.
- You have logged in to the OpenShift Dedicated web console.
- You have logged in to OpenShift Data Science.

Procedure

1. Back up the user's single-user profile.
 - a. Click **Workloads** → **ConfigMaps** in the OpenShift Dedicated web console.
 - b. If it is not already selected, select the **redhat-ods-applications** project from the project list.
 - c. Click the **jupyterhub-singleuser-profile-<username>** ConfigMap.
Replace **<username>** with relevant user name.
 - d. In the **Data** section, click the **Copy** button () to copy the user's data profile to the clipboard.
 - e. Save the contents of the user's data profile to a file.
 - f. Confirm that the file contents are an accurate backup of the user's data profile.
2. Delete the user's persistent volume claim (PVC).
 - a. Click **Storage** → **PersistentVolumeClaims**.
 - b. If it is not already selected, select the **redhat-ods-applications** project from the project list.
 - c. Locate the **jupyterhub-nb-<username>** PVC.
Replace **<username>** with the relevant user name.
 - d. Click the action menu () and select **Delete PersistentVolumeClaim** from the list.

The **Delete PersistentVolumeClaim** dialog appears.

- e. Inspect the dialog and confirm that you are deleting the correct PVC.
 - f. Click **Delete**.
3. Delete the user's ConfigMap.
- a. Click **Workloads** → **ConfigMaps**.
 - b. If it is not already selected, select the **redhat-ods-applications** project from the project list.
 - c. Locate the **jupyterhub-singleuser-profile-<username>** ConfigMap.
Replace **<username>** with the relevant user name.
 - d. Click the action menu (⋮) and select **Delete ConfigMap** from the list.
The **Delete ConfigMap** dialog appears.
 - e. Inspect the dialog and confirm that you are deleting the correct ConfigMap.
 - f. Click **Delete**.

Verification

- The user cannot access JupyterHub any more, and sees a 403 Forbidden error if they try. Note that the user's name remains visible in the JupyterHub administration interface because of a bug in the user deletion process. This is planned for correction in future releases.
- The user's single-user profile, persistent volume claim (PVC), and ConfigMap are not visible in OpenShift Dedicated.

CHAPTER 5. ALLOCATING ADDITIONAL RESOURCES TO OPENSIFT DATA SCIENCE USERS

As a cluster administrator, you can allocate additional resources to a cluster to support compute-intensive data science work. This includes increasing the number of nodes in the cluster and changing the cluster's allocated machine pool.

Prerequisites

- You have an OpenShift Dedicated cluster with an identity provider configured.
- You have credentials for OpenShift Cluster Manager (<https://console.redhat.com/openshift/>).
- You are part of the **cluster-admins** user group in OpenShift Dedicated.
- You have an AWS instance with the capacity to create larger container sizes.
- For compute-intensive operations, you have an AWS instance with enough capacity to accommodate the largest container size, **XL**.

Procedure

1. Log in to OpenShift Cluster Manager (<https://console.redhat.com/openshift/>).
2. Click **Clusters**.
The **Clusters** page opens.
3. Click the name of the cluster you want to allocate additional resources to.
4. Click **Actions** → **Edit node count**.
5. Optional: Select a **Machine pool** from the list.
6. Optional: Select the number of nodes assigned to the machine pool from the **Node count** list.
7. Click **Apply**.

Verification

- The additional resources that you allocated to the cluster are displayed on the **Machine Pools** tab.

CHAPTER 6. MANAGING NOTEBOOK SERVERS

6.1. ACCESSING THE JUPYTERHUB ADMINISTRATION INTERFACE

You can use the JupyterHub administrative interface to control notebook servers in your Red Hat OpenShift Data Science environment.

Prerequisites

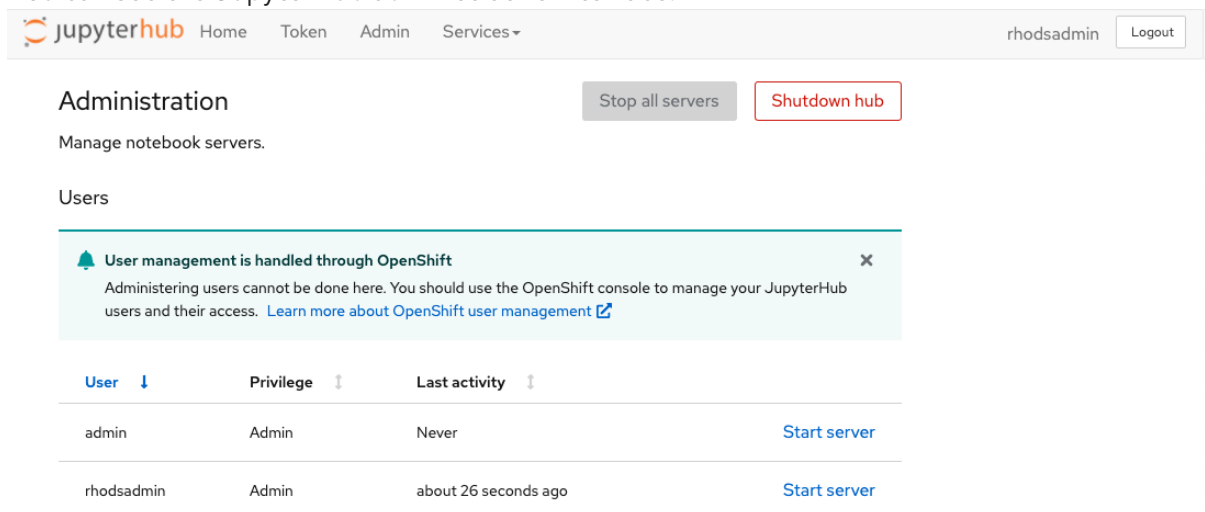
- You are part of the OpenShift Data Science administrator group (by default, **rhods-admins**) in OpenShift Dedicated.

Procedure

- In the OpenShift Data Science interface, click **Enabled**.
- Locate the JupyterHub card and click **Launch**.
- If your notebook server is already running, the JupyterLab interface appears. Click **File** → **Hub Control Panel** to return to JupyterHub.
- Click **Admin** to open the JupyterHub administrative interface.

Verification

- You can see the JupyterHub administrative interface.



The screenshot shows the JupyterHub administrative interface. At the top, there is a navigation bar with 'jupyterhub', 'Home', 'Token', 'Admin', and 'Services'. The user 'rhodsadmin' is logged in, with a 'Logout' button. The main content area is titled 'Administration' and includes buttons for 'Stop all servers' and 'Shutdown hub'. Below this is a 'Users' section with a notification: 'User management is handled through OpenShift. Administering users cannot be done here. You should use the OpenShift console to manage your JupyterHub users and their access. Learn more about OpenShift user management'. Below the notification is a table with columns 'User', 'Privilege', and 'Last activity'.

User	Privilege	Last activity	
admin	Admin	Never	Start server
rhodsadmin	Admin	about 26 seconds ago	Start server

6.2. STARTING NOTEBOOK SERVERS OWNED BY OTHER USERS

Administrators can start a notebook server for another existing user from the JupyterHub administration interface.

Prerequisites

- You are part of the OpenShift Data Science administrator group (by default, **rhods-admins**) in OpenShift Dedicated.

- You have logged in to JupyterHub.

Procedure

- a. In the JupyterHub interface, click the **Admin** tab.
- b. Locate the user whose notebook server you want to start.
- c. Click the **Start server** button.
- d. Fill in the **Start a notebook server** wizard and click **Start server**.
See the *Additional resources* section for help with this step.

Verification

- The JupyterLab home page opens in a new tab.

Additional resources

- [Launching JupyterHub and starting a notebook server](#)
- [Options for notebook server environments](#)

6.3. ACCESSING NOTEBOOK SERVERS OWNED BY OTHER USERS

Administrators can access notebook servers that are owned by other users in order to correct configuration errors or help a data scientist troubleshoot problems with their environment.

Prerequisites

- You are part of the OpenShift Data Science administrator group (by default, **rhods-admins**) in OpenShift Dedicated.
- You have logged in to JupyterHub.
- The notebook server that you want to access is running (started).

Procedure

- a. In the JupyterHub interface, click the **Admin** tab.
- b. Locate the user that the notebook server belongs to.
- c. Click the **Access server** button.

Verification

- The user's notebook server opens in a new tab.

6.4. STOPPING NOTEBOOK SERVERS OWNED BY OTHER USERS

Administrators can stop notebook servers that are owned by other users to reduce resource consumption on the cluster, or as part of removing a user and their resources from the cluster.

Prerequisites

- You are part of the OpenShift Data Science administrator group (by default, **rhods-admins**) in OpenShift Dedicated.
- You have logged in to JupyterHub.
- The notebook server that you want to stop is running (started).

Procedure

1. In the JupyterHub interface, click the **Admin** tab.
2. Stop one or more servers.
 - If you want to stop one or more specific servers:
 - i. Locate the user that the notebook server belongs to.
 - ii. Click the **Stop server** button beside the user.
 - If you want to stop all servers:
 - i. Click the **Stop all** button.
 - ii. Click **OK** to confirm stopping all servers.

Verification

- The **Stop server** button beside each server changes to a **Start server** button when the notebook server has stopped.

CHAPTER 7. BACKING UP STORAGE DATA FROM AMAZON EBS

Red Hat recommends that you back up the data on your persistent volume claims (PVCs) regularly. Backing up your data is particularly important before deleting a user and before uninstalling OpenShift Data Science, as all PVCs are deleted when OpenShift Data Science is uninstalled.

Prerequisites

- You have credentials for OpenShift Cluster Manager (<https://console.redhat.com/openshift/>).
- You have administrator access to the OpenShift Dedicated cluster.
- You have credentials for the Amazon Web Services (AWS) account that the OpenShift Dedicated cluster is deployed under.

Procedure

1. Determine the IDs of the persistent volumes (PVs) that you want to back up.
 - a. In the OpenShift Dedicated web console, change into the **Administrator** perspective.
 - b. Click **Home** → **Projects**.
 - c. Click the **rhods-notebooks** project.
The **Details** page for the project opens.
 - d. Click the **PersistentVolumeClaims** in the **Inventory** section.
The **PersistentVolumeClaims** page opens.
 - e. Note the ID of the persistent volume (PV) that you want to back up.



NOTE

The persistent volumes (PV) that you make a note of are required to identify the correct EBS volume to back up in your AWS instance.

2. Locate the EBS volume containing the PVs that you want to back up.
See [Create Amazon EBS snapshots](#) for more information.
 - a. Log in to AWS (<https://aws.amazon.com>) and ensure that you are viewing the region that your OpenShift Dedicated cluster is deployed in.
 - b. Click **Services**.
 - c. Click **Compute** → **EC2**.
 - d. Click **Elastic Block Storage** → **Volumes** in the side navigation.
The **Volumes** page opens.
 - e. In the search bar, enter the ID of the persistent volume (PV) that you made a note of earlier.
The **Volumes** page reloads to display the search results.
 - f. Click on the volume shown and verify that any **kubernetes.io/created-for/pvc/namespace** tags contain the value **rhods-notebooks**, and any **kubernetes.io/created-for/pvc/name**

tags match the name of the persistent volume that the EC2 volume is being used for, for example, **jupyterhub-nb-user1-pvc**.

3. Back up the EBS volume that contains your persistent volume (PV).
 - a. Right-click on the volume that you want to back up and select **Create Snapshot** from the list.
The **Create Snapshot** page opens.
 - b. Enter a **Description** for the volume.
 - c. Click **Create Snapshot**.
The snapshot of the volume is created.
 - d. Click **Close**.

Verification

- The snapshot that you created is visible on the **Snapshots** page in AWS.

Additional resources

- [Create Amazon EBS snapshots](#)