Red Hat OpenShift AI Cloud Service 1

Managing users

Manage user permissions in Red Hat OpenShift AI

Red Hat OpenShift AI Cloud Service 1 Managing users

Manage user permissions in Red Hat OpenShift AI

Legal Notice

Abstract

Manage user permissions in Red Hat OpenShift AI.

Table of Contents

PREFACE

Users with administrator access to the OpenShift cluster can add, modify, and remove user permissions for Red Hat OpenShift AI.

CHAPTER 1. ADDING USERS

1.1. OVERVIEW OF USER TYPES AND PERMISSIONS

Table 1 describes the Red Hat OpenShift AI user types.

Table 1.1. User types

User Type	Permissions
Data scientists	Data scientists can access and use individual components of Red Hat OpenShift AI, such as Jupyter and data science pipelines. See also Accessing the OpenShift AI dashboard
Administrators	In addition to the actions permitted to a data scientist, administrators can perform these actions: • Configure Red Hat OpenShift AI settings. • Access and manage notebook servers. • Access and manage data science pipeline applications for any data science project. See also OpenShift Dedicated cluster administration or Red Hat OpenShift Service on AWS (ROSA) cluster administration.

By default, all OpenShift users have access to Red Hat OpenShift AI. In addition, users in the OpenShift administrator group (**cluster admins** or **dedicated-admins**), automatically have administrator access in OpenShift AI.

Optionally, if you want to restrict access to your OpenShift AI deployment, you can create specialized user groups for users and administrators.

If you decide to restrict access, and you already have user groups defined in your configured identity provider, you can add these user groups to your OpenShift AI deployment. If you decide to use specialized user groups without adding these groups from an identity provider, you must create the groups in OpenShift and then add users to them.

The user groups configured in OpenShift, **cluster-admins** and **dedicated-admins**, are separate to any specialized user groups for OpenShift AI. There are some operations relevant to OpenShift AI that require the **cluster-admins** or **dedicated-admins** role. Those operations include:

- Adding users to the OpenShift AI user and administrator groups, if you are using specialized groups.

- Removing users from the OpenShift AI user and administrator groups, if you are using specialized groups.

- Managing custom environment and storage configuration for users in OpenShift, such as Jupyter notebook resources, ConfigMaps, and persistent volume claims (PVCs).

- Managing data science pipeline servers and pipeline applications from the OpenShift AI dashboard, or from the command-line interface (CLI), for any data science project.

IMPORTANT

Although users of OpenShift AI and its components are authenticated through OpenShift, session management is separate from authentication. This means that logging out of OpenShift or OpenShift AI does not affect a logged in Jupyter 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.

1.2. DEFINING OPENSHIFT AI ADMINISTRATOR AND USER GROUPS

By default, all users authenticated in OpenShift can access OpenShift AI.

Also by default, users with cluster admin permissions and users in the **dedicated-admins** administrator group are OpenShift AI administrators. A **cluster admin** is a superuser that can perform any action in any project in the OpenShift cluster. When bound to a user with a local binding, they have full control over quota and every action on every resource in the project. The **dedicated-admins** user group applies only to OpenShift Dedicated.

You can define additional administrator and user groups by using the OpenShift AI dashboard.

Prerequisites

- You have logged in to Red Hat OpenShift AI as described in Logging in to OpenShift AI .

- You are part of the administrator group for OpenShift AI in OpenShift.

- The groups that you want to define as administrator and user groups for OpenShift AI already exist in OpenShift.

Procedure

1. From the OpenShift AI dashboard, click **Settings → User management**.

2. Define your OpenShift AI admin groups: Under **Data science administrator groups**, click the text box and select an OpenShift group. Repeat this process to define multiple admin groups.

3. Define your OpenShift AI user groups: Under **Data science user groups**, click the text box and select an OpenShift group. Repeat this process to define multiple user groups.

   IMPORTANT

   The **system:authenticated** setting allows all users authenticated in OpenShift to access OpenShift AI.

4. Click **Save changes**.

Verification

- Administrator users can successfully log in to OpenShift AI and perform administrative functions.

- Non-administrator users can successfully log in to OpenShift AI. They can also access and use individual components, such as Jupyter.

1.3. ADDING USERS TO SPECIALIZED OPENSHIFT AI USER GROUPS

By default, all OpenShift users have access to Red Hat OpenShift AI.

Optionally, you can restrict user access to your OpenShift AI instance by defining specialized user groups. You must grant users permission to access Red Hat OpenShift AI by adding user accounts to the Red Hat OpenShift AI 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 product components in the Red Hat OpenShift AI dashboard, such as data science pipelines, and associated services, such as Jupyter. By default, users in the **user group** have access to data science pipeline applications within data science projects that they created.

The **administrator group** provides the user with access to developer and administrator functions in the Red Hat OpenShift AI dashboard, such as data science pipelines, and associated services, such as Jupyter. Users in the **administrator group** can configure data science pipeline applications in the OpenShift AI dashboard for any data science project.

If you restrict access by using specialized user groups, users that are not in the OpenShift AI user group or administrator group cannot view the dashboard and use associated services, such as Jupyter. They are also unable to access the **Cluster settings** page.

IMPORTANT

If you are using LDAP as your identity provider, you need to configure LDAP syncing to OpenShift. For more information, see Syncing LDAP groups in OpenShift Dedicated or Syncing LDAP groups in Red Hat OpenShift Service on AWS (ROSA)

Follow the steps in this section to add users to your specialized OpenShift AI administrator and user groups.

Note: You can add users in OpenShift AI but you must manage the user lists in the OpenShift web console.

Prerequisites

- You have configured a supported identity provider for your OpenShift cluster.

- You are part of the **cluster-admins** or **dedicated-admins** user group in your OpenShift cluster. The **dedicated-admins** user group applies only to OpenShift Dedicated.

- You have defined an administrator group and user group for OpenShift AI.

Procedure

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

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

   - For administrative users, click the administrator group, for example, **rhoai-admins**.

   - For normal users, click the user group, for example, **rhoai-users**.
     The **Group details** page for that group appears.

3. Click **Actions → Add Users**.

The **Add Users** dialog appears.

4. In the **Users** field, enter the relevant user name to add to the group.

5. Click **Save**.

Verification

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

1.4. VIEWING OPENSHIFT AI USERS

If you have defined specialized user groups for OpenShift AI, you can view the users that belong to these groups.

Prerequisites

- The Red Hat OpenShift AI user group, administrator group, or both exist.

- You have the **cluster-admin** role or you are part of the **dedicated-admins** administrator group. The **dedicated-admins** group applies only to OpenShift Dedicated.

- You have configured a supported identity provider for your OpenShift cluster.

Procedure

1. In the OpenShift 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 the name of your administrator group. for example, **rhoai-admins**.

   - For normal users, click the name of your user group, for example, **rhoai-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 AI.

CHAPTER 2. DELETING USERS

2.1. ABOUT DELETING USERS AND THEIR RESOURCES

If you have administrator access to OpenShift, you can revoke a user's access to Jupyter and delete the user's resources from Red Hat OpenShift AI.

Deleting a user and the user's resources involves the following tasks:

- Before you delete a user from OpenShift AI, it is good practice to back up the data on your persistent volume claims (PVCs).

- Stop notebook servers owned by the user.

- Revoke user access to Jupyter.

- Remove the user from the allowed group in your OpenShift identity provider.

- After you delete a user, delete their associated configuration files from OpenShift.

2.1.1. 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 AI, as all PVCs are deleted when OpenShift AI is uninstalled.

Prerequisites

- You have credentials for Red Hat 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 Amazon Web Services documentation: 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, **jupyter-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

- Amazon Web Services documentation: Create Amazon EBS snapshots

2.1.2. Backing up storage data from Google Persistent Disk

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 AI, as all PVCs are deleted when OpenShift AI is uninstalled.

Prerequisites

- You have credentials for Red Hat OpenShift Cluster Manager
  (https://console.redhat.com/openshift/).

- You have administrator access to the OpenShift Dedicated cluster.

- You have credentials for the Google Cloud Platform (GCP) 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.
      The persistent volume (PV) IDs are required to identify the correct persistent disk to back
      up in your GCP instance.

2. Locate the persistent disk containing the PVs that you want to back up.

   a. Log in to the Google Cloud console (https://console.cloud.google.com) and ensure that
      you are viewing the region that your OpenShift Dedicated cluster is deployed in.

   b. Click the navigation menu (≡) and then click **Compute Engine**.

   c. From the side navigation, under **Storage**, click **Disks**.
      The **Disks** page opens.

   d. In the **Filter** query box, enter the ID of the persistent volume (PV) that you made a note of
      earlier.
      The **Disks** page reloads to display the search results.

   e. Click on the disk 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 persistent disk is being used for, for
      example, **jupyterhub-nb-user1-pvc**.

3. Back up the persistent disk that contains your persistent volume (PV).

   a. Select **CREATE SNAPSHOT** from the top navigation.
      The **Create a snapshot** page opens.

   b. Enter a unique **Name** for the snapshot.

   c. Under **Source disk**, verify the persistent disk you want to back up is displayed.

   d. Change any optional settings as needed.

   e. Click **CREATE**.

The snapshot of the persistent disk is created.

Verification

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

Additional resources

- Google Cloud documentation: Create and manage disk snapshots

2.2. 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

- If you are using specialized OpenShift AI groups, you are part of the administrator group (for example, **rhoai-admins**). If you are not using specialized groups, you are part of the OpenShift Dedicated or Red Hat OpenShift Service on AWS (ROSA) administrator group. For more information, see Adding administrative users.

- You have launched the Jupyter application, as described in Launching Jupyter and starting a notebook server.

- The notebook server that you want to stop is running.

Procedure

1. On the page that opens when you launch Jupyter, click the **Administration** tab.

2. Stop one or more servers.

    - If you want to stop one or more specific servers, perform the following actions:

        i. In the **Users** section, locate the user that the notebook server belongs to.

        ii. To stop the notebook server, perform one of the following actions:

            - Click the action menu (⋮) beside the relevant user and select **Stop server**.

            - Click **View server** beside the relevant user and then click **Stop notebook server**. The **Stop server** dialog box appears.

        iii. Click **Stop server**.

    - If you want to stop all servers, perform the following actions:

        i. Click the **Stop all servers** button.

        ii. Click **OK** to confirm stopping all servers.

Verification

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

2.3. REVOKING USER ACCESS TO JUPYTER

You can revoke a user's access to Jupyter by removing the user from the specialized user groups that define access to OpenShift AI. When you remove a user from the specialized user groups, the user is prevented from accessing the OpenShift AI dashboard and from using associated services that consume resources in your cluster.

> **IMPORTANT**
>
> Follow these steps only if you have implemented specialized user groups to restrict access to OpenShift AI. To completely remove a user from OpenShift AI, 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 **cluster-admins** or **dedicated-admins** user group in your OpenShift cluster. The **dedicated-admins** user group applies only to OpenShift Dedicated.

- You are using specialized user groups for OpenShift AI, and the user is part of the specialized user group, administrator group, or both.

Procedure

1. In the OpenShift 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 the name of your administrator group, for example, **rhoai-admins**.

   - For non-administrator users, click the name of your user group, for example, **rhoai-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 **Workloads → Pods** and ensure that there is no notebook server pod for this user. If you see a pod named **jupyter-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.

- In the OpenShift AI dashboard, check the list of data science projects. Delete any projects that belong to the user.

2.4. CLEANING UP AFTER DELETING USERS

After you remove a user's access to Red Hat OpenShift AI or Jupyter, you must also delete the configuration files for the user from OpenShift. Red Hat recommends that you back up the user's data before removing their configuration files.

Prerequisites

- (Optional) If you want to completely remove the user's access to OpenShift AI, you have removed their credentials from your identity provider.

- You have revoked the user's access to Jupyter.

- You have backed up the user's storage data from Amazon EBS or Google Persistent Disk.

- If you are using specialized OpenShift AI groups, you are part of the administrator group (for example, **rhoai-admins**). If you are not using specialized groups, you are part of the OpenShift Dedicated or Red Hat OpenShift Service on AWS (ROSA) administrator group. For more information, see Adding administrative users.

- You have logged in to the OpenShift web console.

- You have logged in to OpenShift AI.

Procedure

1. Delete the user's persistent volume claim (PVC).

   a. Click **Storage → PersistentVolumeClaims**.

   b. If it is not already selected, select the **rhods-notebooks** project from the project list.

   c. Locate the **jupyter-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**.

2. Delete the user's ConfigMap.

   a. Click **Workloads → ConfigMaps**.

   b. If it is not already selected, select the **rhods-notebooks** 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 Jupyter any more, and sees an "Access permission needed" message if they try.

- The user's single-user profile, persistent volume claim (PVC), and ConfigMap are not visible in OpenShift.