

Cloudera AI

Site Administration

Date published: 2020-07-16

Date modified: 2024-11-21

CLOUDERA

<https://docs.cloudera.com/>

Legal Notice

© Cloudera Inc. 2024. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

Managing Users.....	5
Service Accounts.....	6
Creating a machine user and synchronizing to workspace.....	6
Sync machine users from the Synced team.....	6
Run workloads using a service account.....	7
Configuring Quotas.....	7
Creating Resource profiles.....	8
Disable or deprecate Runtime addons.....	9
Onboarding Business Users.....	11
Adding a Collaborator.....	11
User Roles.....	11
Business Users and Cloudera Machine Learning.....	13
Managing your Personal Account.....	13
Creating a Team.....	13
Managing a Team Account.....	15
Managing a Synced Team.....	15
Monitoring Cloudera Machine Learning Activity.....	16
Tracked User events.....	17
Monitoring User Events.....	20
Monitoring active Models across the Workspace.....	22
Monitoring and alerts.....	23
Application polling endpoint.....	23
Choosing default engine.....	23

Controlling User access to features.....	24
Cloudera Machine Learning email notifications.....	26
Downloading diagnostic bundles for a workspace.....	26
Web session timeouts.....	27
Project garbage collection.....	27
Ephemeral storage.....	27
Installing a non-transparent proxy in a Cloudera Machine Learning environment.....	28
Ports used by Cloudera Machine Learning.....	30
Export Usage List.....	30
Private cluster support.....	30
Enable a private cluster.....	31
User Defined Routing (UDR).....	31
Embed a Cloudera Machine Learning application in an external website.....	33
Setting up Cloudera Machine Learning Workspaces for high volume Workloads.....	34
Host name required by Learning Hub.....	35

Managing Users

This topic describes how to manage an Cloudera Machine Learning Workspace as a site administrator. Site administrators can monitor and manage all user activity across a workspace, add new custom engines, and configure certain security settings.

By default, the first user that logs in to a workspace must always be a site administrator. That is, they must have the MLAdmin role granted by a Cloudera Data Platform PowerUser.



Important: Site administrators have complete access to all activity on the deployment. This includes access to all teams and projects on the deployment, even if they have not been explicitly added as team members or collaborators.

Only site administrators have access to a Site Administration dashboard that can be used to manage the workspace. To access the site administrator dashboard:

1. Go to the Cloudera Machine Learning web application and log in as a site administrator.
2. On the left sidebar, click Site Administration. You will see an array of tabs for all the tasks you can perform as a site administrator.

Release	dev
Domain	
Total Nodes	2
Total Memory	58.10 GiB
Used Memory	22.75 GiB
Total vCPUs	15.25
Used vCPUs	10.19

Monitoring Users

The Users tab on the **Administrator** dashboard displays the complete list of users. You can see which users are currently active, and when a user last logged in to Cloudera Machine Learning. You can search for a user by entering their User ID, Username, or Email in the User quick find box. To modify a user's username, email or permissions, click the Edit button under the **Action** column.



Note: The Disabled checkbox does not have any effect when external authentication is in use.

Synchronizing Users

You can synchronize users within an Cloudera Machine Learning Workspace with those users that have been defined access at the Environment level (through the MLAdmin, MLUser, and MLBusinessUser roles). Doing so for new users enables you to take administrative actions such as setting Team assignments, defining Project Collaborators, and more, all prior to the new users' first time logging in to the Workspace.

To synchronize users, go to **Site Administration Users**, and click **Synchronize Users**. This adds any users defined at the Environment level to the workspace, updates any role changes that have been made, and deactivates any users that have been deactivated.



Note: The Administrator shall periodically perform user synchronization to ensure that users who are deactivated on the environment level are also deactivated in Cloudera Machine Learning.

Synchronizing Groups

Groups of users can be created in the Cloudera Data Platform management console and imported to Cloudera Machine Learning. However, changes made in Cloudera Data Platform do not automatically update in Cloudera Machine Learning. You need to manually trigger an update, using **Sync Teams**. For more information, see *Creating a Team*.

Related Information

[Cloudera Machine Learning email notifications](#)

[Creating a Team](#)

Service Accounts

Service accounts are used by machine users that require a user account, without needing to use an account of an actual user.

Like other users, this machine user can be granted necessary permissions and roles, and be added as a collaborator to projects in order to run workloads. Machine users can also create projects and workloads.

Creating a machine user and synchronizing to workspace

The MLAdmin role is required to create machine users.

Procedure

1. In Management Console, go to User Management.
2. In Actions, click Create Machine User.
3. Enter a name for the machine user and click Create.
4. In the Cloudera Machine Learning Workspaces UI, find your workspace and in Actions, click Manage Access.
5. Search for the machine user name you just created, and in Update Resource Roles, assign the MLWorkspaceAdmin or MLWorkspaceUser role. Click Update Roles.



Note: Machine users can alternatively be assigned to environments.

6. Return to the workspace in Cloudera Machine Learning, and in **Site Administration Users**, click **Run Sync Now** to manually synchronize the users for the workspace.
7. In Site Administration, search for the machine user name.

Sync machine users from the Synced team

You can sync machine users that are part of a synced team to your project.

Procedure

1. In **Management Console User Management Groups**, click **Create Group**.

2. Enter the name for the group, and click Create.
3. Click Add Members to search for and add group members, including machine users.
4. To add the team (group) to your environment, go to Environments Actions Manage Access .
5. Click Update Role to update the role as follows, and click Update Roles.
 - Environment User: Only users who have read access to the environment are synced. Alternatively, you can assign the Environment User role to the machine user.
 - MLAdmin or MLUser role: only users with either role are synced to Cloudera Machine Learning Workspace.
6. Click Synchronize Users and wait for synchronization to complete. Then return to your Cloudera Machine Learning Workspace.
7. In Site Administration Teams , select Sync Teams and then choose the group to sync.
8. Click Create Team, and the team is created in Cloudera Machine Learning.

What to do next

To add members to a synced team, add them in the control plane and synchronize them to Cloudera Machine Learning via the Site Administration Teams Sync Teams option. You cannot add users to a group manually in Cloudera Machine Learning.

Run workloads using a service account

You can run various types of workloads using a service account. First, make sure the service account is available in your project.

1. Create a project, or enter an existing project.
2. In Collaborators, add the service account. Specify the Operator or Admin role and click Add.

Run a job with a service account

1. In Jobs, click New Job.
2. For Run Job as, select Service Account and choose the account from the list.
3. Make other settings as needed, and click Create Job.

Run an application with a service account

1. Click New Application.
2. For Run Job as, select Service Account and choose the account from the list.
3. Make other settings as needed, and click Create Application.

Run a model with a service account

1. In Models, click New Model.
2. For Deploy Model as, select Service Account and choose the account from the list.
3. Make other settings as needed, and click Deploy Model.

Configuring Quotas

This topic describes how to configure CPU, GPU, and memory quotas for users of an Cloudera Machine Learning Workspace.

Before you begin

Required Role: MLAdmin

Make sure you are assigned the MLAdmin role in Cloudera Data Platform. Only users with the MLAdmin role will be logged into Cloudera Machine Learning Workspaces with Site Administrator privileges.

There are two types of quota: Default and Custom. Default quotas apply to all users of the workspace. Custom quotas apply to individual users in the workspace, and take precedence over the default quota.

Procedure

1. Log in to the web interface.
2. Click Cloudera Machine Learning Workspaces, then open the workspace for which you want to set quotas.
3. Click AdminQuotas.
4. Switch the Default Quotas toggle to ON.

This applies a default quota of 2 vCPU and 8 GB memory to each user in the workspace.

If your workspace was provisioned with GPUs, a default quota of 0 GPU per user applies. If you want users to have access to GPUs, you must modify the default quotas as described in the next step.

5. If you want to change the default quotas, click on Default (per user) .
Cloudera Machine Learning displays the Edit default quota dialog box.
6. Enter the CPU, Memory, and GPU quota values that should apply to all users of the workspace.
7. Click Update.
8. To add a custom quota for a specific user, click Add User.
9. Enter the user name, and enter the quotas for CPU, Memory, and GPU.
10. Click Add.

Results

Enabling and modifying quotas will only affect new workloads. If users have already scheduled workloads that exceed the new quota limits, those will continue to run uninterrupted. If a user is over their limit, they will not be able to schedule any more workloads.

What to do next

To specify the maximum number of replicas in a model deployment, go to Site Administration Settings Model Deployment Settings . The default is 9 replicas, and up to 199 can be set.

Creating Resource profiles

Resource profiles define how many vCPUs and how much memory the product will reserve for a particular workload (for example, session, job, model).

About this task

As a site administrator you can create several different vCPU, GPU, and memory configurations which will be available when launching a session/job. When launching a new session, users will be able to select one of the available resource profiles depending on their project's requirements.

Procedure

1. To create resource profiles, go to the Site Administration Runtime/Engine page.
2. Add a new profile under Resource Profiles.

Cloudera recommends that all profiles include at least 2 GB of RAM to avoid out of memory errors for common user operations.

You will see the option to add GPUs to the resource profiles only if your Cloudera Machine Learning hosts are equipped with GPUs, and you have enabled them for use by setting the relevant properties in `cdsw.conf`.

Results

If there are two worker nodes and 10 vCPU available overall, if one user tries to establish a session with 8 vCPU, CDSW will not allow it. The memory and CPU must be contiguous (adjacent to each other). When a user spins a session, the pod triggers on a single node and resources on the same node are utilized. This is expected behavior for Kubernetes.

Figure 1: Resource profiles available when launching a session

The screenshot shows the 'Site Administration / Runtime/Engine' page. On the left, the 'Resource Profiles' table lists available configurations:

Description	vCPU	Memory (GiB)
1 vCPU / 2 GiB Memory	1	2
1 vCPU / 1.75 GiB Memory	1	1.75
1 vCPU, 1.75 GiB memory	1	1.75

On the right, the 'Start A New Session' dialog is open. In the 'Resource Profile' dropdown, the '1 vCPU / 2 GiB Memory' option is selected, matching the profile highlighted in the table on the left.

Disable or deprecate Runtime addons

Disable or deprecate a Spark Runtime addon.

About this task

You can disable or deprecate any Spark Runtime addon from the Runtime/Engine tab of Site Administration.

Procedure

1. Select Site Administration in the left Navigation bar.
2. Select the Runtime/Engine tab.

3. Select Disabled or Deprecated from Actions next to any *Spark* addon.

Site Administration / Runtime/Engine

Runtime Updates

☒ Enable Runtime Updates

New Runtime variants and versions are automatically downloaded and made available on clusters with Internet access. Unch

Hadoop CLI Version Hadoop CLI - CDP 7.2.8 - HOTFIX... ▾

Runtime Addons

Status ▾	Name ▾	ID	Component ▾	Created At
 Available	Hadoop CLI - CDP 7.2.10 - HOTFIX-1 JAVA 8U342	1	HadoopCLI	11/07/2022 12:54 PM
 Available	Hadoop CLI - CDP 7.2.11 - HOTFIX-4 JAVA 8U342	5	HadoopCLI	11/07/2022 12:54 PM
 Available	Hadoop CLI - CDP 7.2.14 - JAVA 8U342	4	HadoopCLI	11/07/2022 12:54 PM
 Available	Hadoop CLI - CDP 7.2.8 - HOTFIX-1 JAVA 8U342	6	HadoopCLI	11/07/2022 12:54 PM
 Deprecated	Spark 2.4.8 - CDE 1.15 - HOTFIX-1	2	Spark	11/07/2022 12:54 PM
 Disabled	Spark 3.2.0 - CDE 1.15 - HOTFIX-2	3	Spark	11/07/2022 12:54 PM



Note: You can also return the status to Available using Actions.

Onboarding Business Users

There are two procedures required for adding Business Users to Cloudera Machine Learning. First, an Administrator ensures the Business User has the correct permissions, and second, a Project Owner adds the Business User as a Collaborator.

Before you begin

Make sure the user is already assigned in your external identity provider, such as LDAP.

About this task

The Admin user performs these steps:

Procedure

1. In Environments, select the correct environment where the Cloudera Machine Learning Workspace is hosted.
2. In Manage Access, search for the user, and add the ML Business User role. Make sure the user does not already have a higher-level permission, such as ML Admin or ML User, either through a direct role assignment or a group membership.
3. Click Update Roles.
4. Inside the Cloudera Machine Learning Workspace, go to **Site Administration > Users**, and click **Synchronize Users**. This adds the necessary Users defined at the Environment level to the workspace, and updates any role changes that have been made.

What to do next

Add the ML Business User as a Collaborator to a Project.

Related Information

[Adding a Collaborator](#)

Adding a Collaborator

Project Owners can add Collaborators to a project.

About this task

The Project Owner performs these steps:

Procedure

1. Go to Collaborators, and enter the user id in the Search box.
2. Choose the User ID, and click Add. The user or team is added with their role displayed.

Results

Now, when the Business User logs in, they are able to access the Applications under this project.

User Roles

Users in Cloudera Machine Learning are assigned one or more of the following roles.



Important: Cloudera Public Cloud allows customers to maintain full ownership and control of their data and workloads and is designed to operate in some of the most restricted Public Cloud environments. Since Cloudera Public Cloud runs in a customer's cloud account, Security and Compliance is a shared responsibility between Cloudera and its Public Cloud customers. User roles form the first layer of security for securing the Cloudera Machine Learning workloads.

It is your responsibility to diligently allocate the permissions to the users. For more information, see *Cloudera's Shared Responsibility Model*.

There are two categories of roles: environment resource roles, which apply to a given Cloudera Data Platform environment, and workspace resource roles, which apply to a single workspace. To use workspace resource roles, you may need to upgrade the workspace or create a new workspace.

If a user has more than one role, then the role with the highest level of permissions takes precedence. If a user is a member of a group, it may gain additional roles through that membership.

Environment resource roles

- **MLAdmin:** Grants a Cloudera Data Platform user the ability to create and delete Cloudera Machine Learning Workspaces within a given Cloudera Data Platform environment. MLAdmins also have Administrator level access to all the workspaces provisioned within this environment. They can run workloads, monitor, and manage all user activity on these workspaces. They can also add the MLUser and MLBusinessUser roles to their assigned environment. This user also needs the account-level role of IAMViewer, in order to access the environment Manage Access page. To create or delete workspaces, this user also needs the EnvironmentAdmin role.
- **MLUser:** Grants a Cloudera Data Platform user the ability to view Cloudera Machine Learning Workspaces provisioned within a given Cloudera Data Platform environment. MLUsers are also able to run workloads on all the workspaces provisioned within this environment.
- **MLBusinessUser:** Grants permission to list Cloudera Machine Learning Workspace for a given Cloudera Data Platform environment. MLBusinessUsers are able to only view applications deployed under the projects that they have been added to as a Business User.

Workspace resource roles

Workspace roles are for users who are granted access to only a single workspace.

- **MLWorkspaceAdmin:** Grants permission to manage all Cloudera Machine Learning workloads and settings inside a specific workspace. To perform resource role assignment, the IAMViewer role is also needed. A user with this role can administer the workspace, but is not able to utilize Cloudera Data Platform APIs that modify a workspace (for example, creating or upgrading workspaces).
- **MLWorkspaceBusinessUser:** Grants permission to view shared Cloudera Machine Learning applications inside a specific workspace.
- **MLWorkspaceUser:** Grants permission to run Cloudera Machine Learning workloads inside a specific workspace.

Using the workspace resource roles

A power user or account administrator must assign the first MLWorkspaceAdmin to a workspace. Subsequently, if the MLWorkspaceAdmin also has the IAMViewer role, they can assign resource roles to the workspace.

An MLAdmin (an environment resource role) is not automatically able assign workspace resource roles on the Manage access page. A role such as MLWorkspaceAdmin is needed to do this.

You can check the permissions for a given resource role by clicking the Information icon by each resource role shown in User Management, on the Resources tab for a user, or in a Cloudera Data Platform user profile.



Note: Any user that lists users or assigns resource roles also needs the account-level role of IAMViewer.

Business Users and Cloudera Machine Learning

A user is treated as a Business User inside of Cloudera Machine Learning if they are granted the `MLBusinessUser` role on the Environment of the given Cloudera Machine Learning Workspace. Inside the workspace, a Business User is able to access and view applications, but does not have privileges to access any other workloads in the workspace.

Logging in as a Business User

When you log in as a Business User, the only page you see is the Applications page. The page shows any applications associated with any projects that you have been added to as a Collaborator, even though you do not have rights to access the other assets associated with those projects.

In order for applications to appear in your view, contact the Project Owner to add you as a Collaborator to the project. If you have not been added to any projects, or none of the projects that you have been added to have applications, the Applications page displays the message, You currently don't have any applications.

Managing your Personal Account

You can edit personal account settings such as email, SSH keys and Hadoop credentials.

About this task

You can also access your personal account settings by clicking Account settings in the upper right-hand corner drop-down menu. This option will always take you to your personal settings page, irrespective of the context you are currently in.

Procedure

1. Sign in to Cloudera Machine Learning.
2. From the upper right drop-down menu, switch context to your personal account.
3. Click Settings.

Profile

You can modify your name, email, and bio on this page.

Teams

This page lists the teams you are a part of and the role assigned to you for each team.

SSH Keys

Your public SSH key resides here. SSH keys provide a useful way to access to external resources such as databases or remote Git repositories. For instructions, see *SSH Keys*.

Related Information

[SSH Keys](#)

Creating a Team

Users who work together on more than one project and want to facilitate collaboration can create a Team. Teams enable you to efficiently manage the users assigned to projects.

About this task

Team projects are owned by the team, rather than an individual user. Team administrators can add or remove members at any time, assigning each member different permissions. A team cannot be deleted and at least one member must be there in the team.

Site Administration

Overview Users **Teams** Usage Quotas Models Runtime Data Connections Security AMPs Learning Hub Settings Support

Create Team

* Name

Description

Team Type
☐ Local ☒ Synced Team

Add Groups

Select Group	Select Role	Add
Name	Viewer	
	Operator	
	Contributor	
	Admin	

No data

Cancel Create Team

Procedure

1. In Site Administration Teams , select New Team.

2. Enter the name of the team.

3. Select Local or Synced Team.

Cloudera Data Platform manages the member data of a Synced Team. The members and information about the members of a Local team is not managed by Cloudera Data Platform.

4. If Synced Team is selected, under Add Groups, select a group name and the role for the group and click Add. You can add multiple groups and roles using the Add option.



Note: By default, each member will inherit the role of the groups they belong to. If a member belongs to multiple groups, their effective role in the team is the highest role assigned to the member (Viewer < Operator < Contributor < Admin).

5. Enter a Description, if needed.

6. Add or invite team members. Team members can have one of the following privilege levels:

- Viewer - The Viewer has read-only access to team projects. The Viewers cannot create new projects within the team but can be added to existing ones.
- Operator - The Operator has read-only access to team projects. Additionally, Operators can start and stop existing jobs in the projects that they have access to.
- Contributor - The Contributor has write-level access to all team projects to all team projects with Team or Public visibility. The Contributor can create new projects within the team. They can also be added to existing team projects.
- Admin - The Administrator has complete access to all team projects, can add new team members, and modify team account information. The creator of the team is assigned the Administrator privilege, and can also assign other team members the Administrator privilege. Each team must have at least one Administrator user.

7. Select Create Team.

8. Select Sync Teams to update the teams with information in the Management Console.

Managing a Team Account

Team administrators can modify account information, add or invite new team members, and view/edit privileges of existing members.

Procedure

1. From the upper right drop-down menu, switch context to the team account.
2. Click Settings to open up the Account Settings dashboard.
3. Modify any of the following settings:

Profile

Modify the team description on this page.

Members

You can add new team members on this page, and modify privilege levels for existing members.

SSH Keys

The team's public SSH key resides here. Team SSH keys provide a useful way to give an entire team access to external resources such as databases. For instructions, see *SSH Keys*. Generally, team SSH keys should not be used to authenticate against Git repositories. Use your personal key instead.

Related Information

[SSH Keys](#)

Managing a Synced Team

Team administrators and Site administrators can view members of a group, delete a group within a team, update roles for a group within a team, and update a custom role for a member within a group.

Viewing members of a group

You can view the members of a group along with their roles for a particular group within a team.

1. In the Cloudera Data Platform console, click the Cloudera Machine Learning tile.
The **Home** page displays.
2. Click **Site Administration** in the left navigation menu.
3. Click Teams tab.
4. In the Teams page, click on the group name to view the members' information.
5. In the Groups tab, all the groups and their role is displayed.
6. Click the Members tab to view the list of members, their highest role, and all the groups that they belong to.

The Inherit label is displayed if the role is inherited as part of the group. The Custom Role Set label is displayed if a custom role is assigned to the user.

Adding a custom role for a member

By default, members inherit the role of the group. You can set a custom role for a specific member within a group.

1. In the Cloudera Data Platform console, click the Cloudera Machine Learning tile.
The **Home** page displays.
2. Click **Site Administration** in the left navigation menu.
3. Click Teams tab.
4. In the Teams page, click on the group name.

5. Click the Members tab.

List of the members, their highest role, and all the groups that they belong to is displayed.

6. Under Role, select the role you want to assign to the member from the drop-down list.

The Custom Role Set label is displayed after a custom role is assigned to the user.



Note: When you change the role of the group, the custom role of the member does not change. You must change the custom role to the Inherit role from the Role drop-down list for the member to inherit the group's role.

Updating the role of a group within a team

You can update the role for a particular group within a team.

1. In the Cloudera Data Platform console, click the Cloudera Machine Learning tile.

The **Home** page displays.

2. Click **Site Administration** in the left navigation menu.
3. Click Teams tab.
4. In the Teams page, click on the group name.
5. In the Groups tab, under Role, select the role you want to assign to the group from the drop-down list.

The new role will be implemented for all members who inherit their roles within the group. Members with custom roles will remain unaffected.


Monitoring Cloudera Machine Learning Activity

This topic describes how to monitor user activity on an Cloudera Machine Learning Workspace.

Required Role: Site Administrator

The **Admin Overview** tab displays basic information about your deployment, such as the number of users signed up, the number of teams and projects created, memory used, and some average job scheduling and run times. You can also see the version of Cloudera Machine Learning you are currently running.

The **Admin Activity** tab of the dashboard displays the following time series charts. These graphs should help site administrators identify basic usage patterns, understand how cluster resources are being utilized over time, and how they are being distributed among teams and users.



Important: The graphs and numbers on the Admin Activity page do not account for any resources used by active models on the deployment. For that information, go to Admin Models page.

- CPU - Total number of CPUs requested by sessions running at this time.

Note that code running inside an n-CPU session, job, experiment or model replica can access at least n CPUs worth of CPU time. Each user pod can utilize all of its host's CPU resources except the amount requested by other user workloads or Cloudera Machine Learning application components. For example, a 1-core Python session can use more than 1 core if other cores have not been requested by other user workloads or Cloudera Machine Learning application components.

- Memory - Total memory (in GiB) requested by sessions running at this time.
- GPU - Total number of GPUs requested by sessions running at this time.
- Runs - Total number of sessions and jobs running at this time.
- Lag - Depicts session scheduling and startup times.
 - Scheduling Duration: The amount of time it took for a session pod to be scheduled on the cluster.
 - Starting Duration: The amount of time it took for a session to be ready for user input. This is the amount of time since a pod was scheduled on the cluster until code could be executed.

The Export Sessions List provides a CSV export file of the columns listed in the table. It is important to note that the exported duration column is in seconds for a more detailed output.

Tracked User events

The tables on this page describe the user events that are logged by Cloudera Machine Learning.

Table 1: Database Columns

When you query the user_events table, the following information can be returned:

Information	Description
id	The ID assigned to the event.
user_id	The UUID of the user who triggered the event.

Information	Description
ipaddr	The IP address of the user or component that triggered the event. 127.0.0.1 indicates an internal component.
user agent	The user agent for this action, such as the web browser. For example: Mozilla/5.0 (X11; Linux x86_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/51.0.2704.103 Safari/537.36
event_name	The event that was logged. The tables on this page list possible events.
description	This field contains the model name and ID, the user type (NORMAL or ADMIN), and the username.
created_at	The date (YYYY-MM-DD format) and time (24-hour clock) the event occurred .

Table 2: Events Related to Engines

Event	Description
engine environment vars updated	-
engine mount created	-
engine mount deleted	-
engine mount updated	-
engine profile created	-
engine profile deleted	-
engine profile updated	-

Table 3: Events Related to Experiments

Event	Description
experiment run created	-
experiment run repeated	-
experiment run cancelled	-

Table 4: Events Related to Files

Event	Description
file downloaded	-
file updated	-
file deleted	-
file copied	-
file renamed	-
file linked	The logged event indicates when a symlink is created for a file or directory.
directory uploaded	-

Table 5: Events Related to Models

Event	Description
model created	-
model deleted	-

Table 6: Events Related to Jobs

Event	Description
job created	-
job started	-
stopped all runs for job	-
job shared with user	-
job unshared with user	-
job sharing updated	<p>The logged event indicates when the sharing status for a job is changed from one of the following options to another:</p> <ul style="list-style-type: none"> • All anonymous users with the link • All authenticated users with the link • Specific users and teams

Table 7: Events Related to Licenses

Event	Description
license created	-
license deleted	-

Table 8: Events Related to Projects

Event	Description
project created	-
project updated	-
project deleted	-
collaborator added	-
collaborator removed	-
collaborator invited	-

Table 9: Events Related to Sessions

Event	Description
session launched	-
session terminated	-
session stopped	-
session shared with user	-
session unshared with user	-

Event	Description
update session sharing status	<p>The logged event indicates when the sharing status for a session is changed from one of the following options to another:</p> <ul style="list-style-type: none"> All anonymous users with the link All authenticated users with the link Specific users and teams

Table 10: Events Related to Admin Settings

Event	Description
site config updated	The logged event indicates when a setting on the Admin Settings page is changed.

Table 11: Events Related to Teams

Event	Description
add member to team	-
delete team member	-
update team member	-

Table 12: Events Related to Users

Event	Description
forgot password	-
password reset	-
update user	If the logged event shows that a user is banned, that means that the user account has been deactivated and does not count toward the license.
user signup	-
user login	The logged event includes the authorization method, LDAP/SAML or local.
user logout	-
ldap/saml user creation	The logged event indicates when a user is created with LDAP or SAML.

Monitoring User Events

This topic shows you how to query the PostgreSQL database that is embedded within the Cloudera Machine Learning deployment to monitor or audit user events.

About this task

Querying the PostgreSQL database that is embedded within the Cloudera Machine Learning deployment requires root access to the Cloudera Machine Learning Master host.

Procedure

1. SSH to the Cloudera Machine Learning Master host and log in as root.

For example, the following command connects to `cdsw-master-host` as root:

```
ssh root@cdsw-master-host.yourcdswdomain.com
```

2. Get the name of the database pod:

```
kubectl get pods -l role=db
```

The command returns information similar to the following example:

NAME	READY	STATUS	RESTARTS	AGE
db-86bbb69b54-d5q88	1/1	Running	0	4h46m

3. Enter the following command to log into the database as the sense user:

```
kubectl exec <database pod> -ti -- psql -U sense
```

For example, the following command logs in to the database on pod db-86bbb69b54-d5q88:

```
kubectl exec db-86bbb69b54-d5q88 -ti -- psql -U sense
```

You are logged into the database as the sense user.

4. Run queries against the user_events table.

For example, run the following query to view the most recent user event:

```
select * from user_events order by created_at DESC LIMIT 1
```

The command returns information similar to the following:

id	3658
user_id	273
ipaddr	::ffff:127.0.0.1
user_agent	node-superagent/2.3.0
event_name	model created
description	{"model":"Simple Model 1559154287-ex5yn","modelId":"50","userType":"NORMAL","username":"DonaldBatz"}
created_at	2019-05-29 18:24:47.65449

5. Optionally, you can export the user events to a CSV file for further analysis:

- a) Copy the user_events table to a CSV file:

```
copy user_events to '/tmp/user_events.csv' DELIMITER ',' CSV HEADER;
```

- b) Find the container that the database runs on:

```
docker ps | grep db-86bbb
```

The command returns output similar to the following:

```
c56d04bbd58 c230b2f564da "docker-entrypoint..." 7 days ago Up 7 days k8s_db-db-86bbb69b54-fcfm6_default_8b2dd23d-88b9-11e9-bc34-0245eb679f96_0
```

The first entry is the container ID.

- c) Copy the user_events.csv file out of the container into a temporary directory on the Master host:

```
docker cp <container ID>:/tmp/user_events.csv /tmp/user_events.csv
```

For example:

```
docker cp 8c56d04bbd58:/tmp/user_events.csv /tmp/user_events.csv
```

- d) Use SCP to copy /tmp/user_events.csv from the Cloudera Machine Learning Master host to a destination of your choice.

For example, run the following command on your local machine to copy user_events.csv to a local directory named events:

```
scp root@cdsw-master-host.yourcdswdomain.com:/tmp/user_events.csv /local/directory/events/
```

What to do next

For information about the different user events, see *Tracked User Events*.

Related Information

[Tracked User events](#)

Monitoring active Models across the Workspace

This topic describes how to monitor all active models currently deployed on your workspace.


What is an active Model?

A model that is in the Deploying, Deployed, or Stopping stages is referred to as an active model.

Monitoring all active Models across the Workspace

Required Role: Site Administrator

To see a complete list of all the models that have been deployed on a deployment, and review resource usage across the deployment by models alone, go to **Admin Models**. On this page, site administrators can also Stop/Restart/Rebuild any of the currently deployed models.



Monitoring and alerts

Cloudera Machine Learning leverages Cloudera Data Platform Monitoring based on Prometheus and Grafana to provide dashboards that allow you to monitor how CPU, memory, storage, and other resources are being consumed by your Cloudera Machine Learning Workspaces.

Prometheus is an internal data source that is auto-populated with resource consumption data for each deployment. Grafana is the monitoring dashboard that allows you to create visualizations for resource consumption data from Prometheus. By default, Cloudera Machine Learning provides three Grafana dashboards: K8 Cluster, K8s Containers, and K8s Node. You can extend these dashboards or create more panels for other metrics. For more information, see the Grafana documentation.

Related Information

[Grafana documentation](#)

Application polling endpoint

The Cloudera Machine Learning server periodically polls applications for their status. The default polling endpoint is the root endpoint (/), but a custom polling endpoint can be specified if the server or other application has difficulty with the default endpoint.

When creating or modifying an application, you can specify a new value for the CDSW_APP_POLLING_ENDPOINT environmental variable. Just replace the default value / that is shown. For more information, see *Analytical Applications*.

You can also set the environmental value in *Project Settings Advanced* . In this case, any setting made here can be overridden by settings in a given application. However, settings made in *Project Settings Advanced* also apply when polling sessions.

Related Information

[Analytical Applications](#)

Choosing default engine

This topic describes how to choose a default engine for creating projects.

Before you begin

Required Role: MLAdmin



Note: On Private Cloud, the corresponding role is EnvironmentAdmin.

Make sure you are assigned the MLAdmin role in Cloudera Data Platform. Only users with the MLAdmin role will be logged into Cloudera Machine Learning Workspaces with Site Administrator privileges.

There are two types of default engines: and Legacy Engines. However, legacy engines are deprecated in the current release and project settings default to ML Runtime.

Legacy engines contain the machinery necessary to run sessions using all four interpreter options that Cloudera Machine Learning currently supports (Python 2, Python 3, R and Scala) and other support utilities (C and Fortran compilers, LaTeX, etc.). ML Runtimes are thinner and more lightweight than legacy engines. Rather than supporting multiple programming languages in a single engine, each Runtime variant supports a single interpreter version and a subset of utilities and libraries to run the user's code in Sessions, Jobs, Experiments, Models, or Applications.

Procedure

1. Log in to the web interface.
2. Click Cloudera Machine Learning Workspaces , then open the workspace for which you want to set Default Engine.
3. Click Admin Runtime/Engine .
4. Choose the Default Engine you would like to use as the default for all newly created projects in this workspace.



Note: Legacy Engines are deprecated in this release and Cloudera recommends using Runtime.

5. Modify the remaining information on the page:
 - Resource Profiles listed in the table are selectable resource options for both legacy Engines and ML Runtime (for example, when starting a Session or Job)
 - The remaining information on the page applies to site-level settings specific for legacy Engines.

Controlling User access to features

Cloudera Machine Learning provides Site Administrators with the ability to restrict or hide specific functionality that non-Site Administrator users have access to in the UI. For example, a site administrator can hide the models and experiments features from the Cloudera Machine Learning Workspace UI.

The settings on this page can be configured through the Security and Settings tabs on the Administration page.

Table 13: Security Tab

Property	Description
Allow remote editing	Disable this property to prevent users from connecting to the Cloudera Machine Learning deployment with cdswctl and using local IDEs, such as PyCharm.

Property	Description
Allow only session creators to run commands on active sessions	By default, a user's permission to active sessions in a project is the same as the user's permission to that project, which is determined by the combination of the user's permission as a project collaborator, the user's permission in the team if this is a team project, and whether the user is a Site Administrator. By checking this checkbox, only the user that created the active session will be able to run commands in that session. No other users, regardless of their permissions in the team or as project collaborators, will be able to run commands on active sessions that are not created by them. Even Site Administrators will not be able to run commands in other users' active sessions.
Allow console output sharing	Disable this property to remove the Share button from the project workspace and workbench UI as well as disable access to all shared console outputs across the deployment. Note that re-enabling this property does not automatically grant access to previously shared consoles. You will need to manually share each console again
Allow anonymous access to shared console outputs	Disable this property to require users to be logged in to access shared console outputs.
Allow file upload/download through UI	Use this checkbox to show/hide file upload/download UI in the project workspace. When disabled, Cloudera Machine Learning API will forbid request of downloading file(s) as attachment. Note that the backend API to upload/edit/read the project files are intact regardless of this change in order to support basic Cloudera Machine Learning functionality such as file edit/read.

Table 14: Settings Tab

Property	Description
Require invitation to sign up	Enable this property to send email invitations to users when you add them to a group. To send email, an SMTP server must first be configured in Settings Email .
Allow users to create public projects	Disable this property to restrict users from creating new public projects. Site Administrators will have to create any new public projects.
Allow Legacy Engine users to use the Python 2 kernel	Enable this property to allow Legacy Engine users to select the Python 2 kernel when creating a job. Python 2 is disabled by default.
Allow users to create projects	Disable this property to restrict users from creating new projects. Site Administrators will have to create any new projects.
Allow users to create teams	Disable this property to restrict users from creating new teams. Site Administrators will have to create any new teams.
Allow users to run experiments	Disable this property to hide the Experiments feature in the UI. Note that this property does not affect any active experiments. It will also not stop any experiments that have already been queued for execution.
Allow users to create models	Disable this property to hide the Models feature in the UI. Note that this property does not affect any active models. In particular, if you do not stop active models before hiding the Models feature, they continue to serve requests and consume computing resources in the background.
Allow users to create applications	Disable this property to hide the Applications feature in the UI. Note that this property does not affect any active applications. In particular, if you do not stop active applications before hiding the feature, they continue to serve requests and consume computing resources in the background.

Cloudera Machine Learning email notifications

Cloudera Machine Learning allows you to send email notifications when you add collaborators to a project, share a project with a colleague, and for job status updates (email recipients are configured per-job). This topic shows you how to specify email address for such outbound communications.

Note that email notifications are not currently enabled by default. Emails are not sent when you create a new project. Email preferences cannot currently be configured at an individual user level.

Option 1: If your existing corporate SMTP server is accessible from the VPC where your Cloudera Machine Learning Workspace is running, you can continue to use that server. Go to the **Admin Settings** tab to specify an email address for outbound invitations and job notifications.

Option 2: If your existing SMTP solution cannot be used, consider using an email service provided by your cloud provider service. For example, Amazon provides Amazon Simple Email Service (Amazon SES).

Mail relay hosts often limit the authenticated sender reply address. Make sure to select a No reply email which you are allowed to use, otherwise email sending may fail.

Downloading diagnostic bundles for a workspace

This topic describes how to download diagnostic bundles for an Cloudera Machine Learning Workspace.

Before you begin

Required Role: MLAdmin

Make sure you are assigned the MLAdmin role in Cloudera Data Platform. Only users with the MLAdmin role will be logged into Cloudera Machine Learning Workspaces with Site Administrator privileges.

Procedure

1. Log in to the web interface.
2. In Cloudera Machine Learning Workspaces, select a workspace.
3. Select **Site Administration Support Generate Log Archive**.
4. Select the time period from the dropdown.
5. Ensure **Include Engines** is selected if engine logs are needed (included by default).
6. Select **Send to Cloudera** to send the diagnostic logs to Cloudera Support.
7. Select **Create** to generate the logs.
8. When Status is Complete, select **Download** to download the diagnostics bundles to your machine.

What to do next

The data in the contained bundles may be incomplete. If it does not contain logs for time period you are looking for, there are a number of possible reasons:

- There is a delay between the time the logs are initially generated by a workload and the time they are visible in cloud storage. This may be approximately 1 minute due to buffering during streaming, but can be significantly longer due to eventual consistency in the cloud storage.
- Another user or process may have deleted data from your bucket; this is beyond the control of Cloudera Machine Learning.
- There may be a misconfiguration or an invalid parameter in your request. Retrieving logs requires a valid cloud storage location to be configured for logging, as well as authentication for Cloudera Machine Learning to be set up properly for it. Requests must pertain to a valid engine in a valid project.

Web session timeouts

You can set web sessions to time out and require the user to log in again. This time limit is not based on activity, it is the maximum time allowed for a web session.

You can set timeout limits for Users and Admin Users in `Site AdministrationSecurity`.

- **User Web Browser Timeout (minutes)** - This timeout sets the default maximum length of time that a web browser session can remain inactive. You remain logged in if you are actively using the session. If you are not active, then after a 5-minute warning, you are automatically logged out. Any changes to the setting take effect for any subsequent user logins.
- **Admin User Web Browser Timeout (minutes)** - This timeout sets the default maximum length of time that a web browser session for an Admin user can remain inactive. You remain logged in if you are actively using the session. If you are not active, then after a 5-minute warning, you are automatically logged out. Any changes to the setting take effect for any subsequent Admin user logins.

Project garbage collection

Marks orphaned files for deletion from a project and cleans up projects that are marked for deletion.

Procedure

1. Click `Site Administration Settings`.
2. Scroll to `Project Garbage Collection`.

Click `Garbage Collect Projects` to permanently delete projects marked for deletion.

Click `Clean Up Orphaned Projects` to mark orphaned projects for deletion.

Results

Orphaned project files are marked for deletion. All files marked for deletion are permanently deleted when you click `Garbage Collect Projects`.

Ephemeral storage

Ephemeral storage space is scratch space a Cloudera Machine Learning session, job, application or model can use. This feature helps in better scheduling of Cloudera Machine Learning pods, and provides a safety valve to ensure runaway computations do not consume all available scratch space on the node.

By default, each user pod in Cloudera Machine Learning is allocated 0 GB of scratch space, and it is allowed to use up to 10 GB. These settings can be applied to an entire site, or on a per-project basis.

Change site-wide ephemeral storage configuration

In `Site Administration Settings Advanced`, you can see the fields to change the ephemeral storage request (minimum) and maximum limit.

Ephemeral Storage Settings**Ephemeral Storage Request (in GB)**

The amount of scratch space requested by the session pod.

Ephemeral Storage Limit (in GB)

The maximum amount of scratch space the session pod is permitted to use. Kubernetes terminates the pod if it exceeds this limit.

[Update](#)**Override Site-wide ephemeral storage configuration**

If you want to customize the ephemeral storage settings, you can do so on a per-project basis. Open your project, then click on **Project Settings Advanced** and adjust the ephemeral storage parameters.

Ephemeral Storage Settings

The amount of scratch space requested by the session pod. The value set here is for the specific project.

Ephemeral Storage Request

GB

The maximum amount of scratch space the session pod is permitted to use. Kubernetes terminates the pod if it exceeds this limit. The value set here is for the specific project.

Ephemeral Storage Limit

GB

[Apply](#)

Click on the below button to reset the project-level ephemeral storage values to match the values set on site level.

[Reset Ephemeral Storage](#)**AWS Known Issues**

There is a known issue with the cluster autoscaler that affects autoscaling from 0->1 if a non-zero value for Ephemeral Storage Request is set. This affects both CPU and GPU node groups of the Cloudera Machine Learning Workspace. The autoscaler throws the following error when this happens:

```
pod didn't trigger scale-up: 1 Insufficient ephemeral-storage
```

This is occurring even though the nodes in the Cloudera Machine Learning autoscaling groups have sufficient ephemeral storage space in their group templates. See this [github issue](#) for details. Even though the issue is closed, the problem still persists.

The issue only affects node groups that have [0, x] autoscaling range.

Set the Ephemeral Storage Request value to 0 in both the site-wide and project settings if you run into this issue.

Installing a non-transparent proxy in a Cloudera Machine Learning environment



If Cloudera Machine Learning is used in an air-gapped environment, a proxy configuration is not mandatory. If a non-transparent proxy is used, then certain endpoints must be added to the allowed list for the proxy.

If your Cloudera Data Platform Private Cloud deployment uses a non-transparent network proxy, configure proxy hosts that the workloads can use for connections with CML workspaces. You can configure the proxy configuration values from the Management Console.



Note: The settings configured using this procedure reflect in newly provisioned CML workspaces in a CDP Private Cloud Experiences deployment using the Experiences Compute Service (ECS). In an OpenShift deployment, the default values are used.

1. Sign in to the CDP console.
2. Click Management Console.
3. On the Management Console home page, select **Administration Networks** to view the Networks page.
4. Configure the following options for the proxy values:

Field	Description
HTTPS Proxy	<p>The HTTP or HTTPS proxy connection string for use in connections with CML workspaces. You must specify this connection string in the form: <code>http(s)://<username>:<password>@<host>:<port></code>.</p> <p> Note: The <code><username></code> and <code><password></code> parameters are optional. You can specify the connection proxy string without these parameters.</p>
HTTP Proxy	<p>The HTTP or HTTPS proxy connection string for use in connections with CML workspaces. You must specify this connection string in the form: <code>http(s)://<username>:<password>@<host>:<port></code>.</p> <p> Note: The <code><username></code> and <code><password></code> parameters are optional. You can specify the connection proxy string without these parameters.</p>
No Proxy	<p>Comma-separated list of hostnames, IP addresses, or hostnames and IP addresses that should not be accessed through the specified HTTPS or HTTP proxy URLs.</p> <p>In case of ECS deployments, you must include no-proxy URLs for the following:</p> <ul style="list-style-type: none"> • All the ECS hosts in your deployment • Any CDP Private Cloud Base cluster that you want to access • CIDR IP addresses for internal operations in the ECS cluster: 10.42.0.0/16 and 10.43.0.0/16

5. Click Save
6. Ensure that the following endpoint is allowed:

Description/Usage	CDP service	Destination	Protocol and Authentication	IP Protocol/Port	Comments
Accelerators for ML Projects (AMPs)	Machine Learning	<code>https://raw.githubusercontent.com</code>	HTTPS	TCP/443	Files for AMPs are hosted on GitHub.

Description/ Usage	CDP service	Destination	Protocol and Authentication	IP Protocol/ Port	Comments
		https://github.com			

Ports used by Cloudera Machine Learning

Certain ports must be accessible through the firewall for proper operation of Cloudera Machine Learning.

Firewall restrictions must be disabled across Cloudera Machine Learning and Cloudera Data Platform cluster hosts. Internally, the Cloudera Machine Learning master and worker hosts require full connectivity with no firewalls. Externally, end users connect to Cloudera Machine Learning exclusively through a web server running on the master host, and therefore do not need direct access to any other internal Cloudera Machine Learning or Cloudera Data Platform services.

Communication with the Cloudera Data Platform cluster	Cloudera Machine Learning -> Cloudera Data Platform As a gateway service, Cloudera Machine Learning must have access to all the ports used by Cloudera Data Platform and Cloudera Manager .
Communication with the Web Browser	The Cloudera Machine Learning web application is available at port 80. HTTPS access is available over port 443.

Export Usage List

You can export a list of sessions, jobs, workers, and experiments. You can either download a complete list of workloads or you can filter the workloads by date to download a more concise list. Timestamps in the list are given in Coordinated Universal Time (UTC).

Procedure

1. Select Site Administration in the navigation pane.
2. Select the Usage tab.
3. If you want a list of workloads specific to a date range, you can filter the list of workloads by setting the Date Range.
4. Select Export Usage List to download the list of workloads.

Results

The list downloads to your computer as a .csv file.

Private cluster support

Each type of network architecture supported by Cloudera Data Platform has a unique set of tradeoffs among ease of setup, security, workloads (Experiences) supported, and so on.

Private Cloud Clusters provide a simple way to create a secure cluster, where the API server and the workloads themselves only rely on private IP addresses that are not accessible from the internet. Connectivity to the cluster from the Cloudera Data Platform control plane is provided by the Cluster Connectivity Manager v2 (CCM v2). CCMv2

uses an agent running in the cluster, and an inverting proxy running on Cloudera Data Platform, which creates a HTTPS tunnel between the workload and the control plane.



Note: Private Cloud clusters are Generally Available (GA) on AWS, but are in Preview on Azure, and require an entitlement.

Enable a private cluster

To enable a private cluster, select the option when provisioning the workspace.

Procedure

1. In Cloudera Machine Learning Workspaces, select Provision Workspace.
2. Enter a Workspace Name, and select Environment.
3. Select the Advanced Options toggle.
4. In Network Settings, select Enable Fully Private Cluster.
5. Make any other settings needed, and select Provision Workspace.

Network Settings

Subnets ⓘ

Select Subnets

Load Balancer Source Ranges ⓘ

0.0.0.0/0

☐ Enable Fully Private Cluster

☐ Enable Public IP Address for Load Balancer

Results

The workspace is provisioned using a fully private cluster.

What to do next

User Defined Routing (UDR)

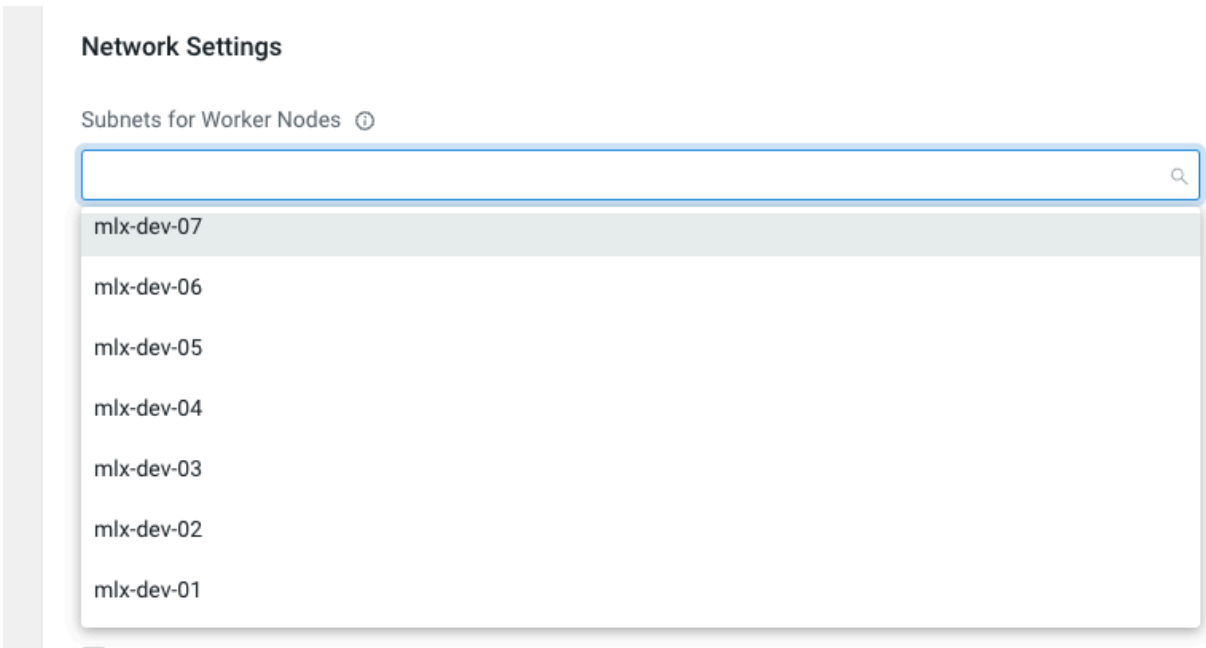
With the Fully Private Cluster configuration, Azure still creates some public IP resources to support load balancer egress. If necessary, you can avoid creating public IP addresses in the Cloudera Machine Learning cluster by using a User Defined Routing (UDR) table. A UDR table can be configured in the cluster subnet to route packets to a customer-configured firewall, for example to limit internet access or analyze traffic. For more information on setting up UDR, see the Microsoft articles [Virtual appliance scenario](#) or [Virtual network traffic routing](#).

About this task

To utilize a UDR and firewall in the Azure Cloudera Machine Learning Private Cloud cluster, select the following when setting up the cluster.

Procedure

- 1. Select a subnet with a default route configuration to forward the traffic to the network appliance or firewall.



- 2. Create load balancers with private IP addresses. This is the default choice when creating clusters in Cloudera Machine Learning.

Network Settings

Subnets for Worker Nodes ⓘ

mlx-dev-01 ▼

Load Balancer Source Ranges ⓘ

0.0.0.0/0 - +

☒ Enable Fully Private Cluster

3. Select Enable User Defined Routing.

Network Settings

Subnets for Worker Nodes ⓘ

mlx-dev-01

Load Balancer Source Ranges ⓘ

0.0.0.0/0

☒ When enabled, the Azure CML Workspace will use the configured UDR.

☒ Enable User Defined Routing ⓘ

Embed a Cloudera Machine Learning application in an external website

You can embed a Cloudera Machine Learning application into an I-frame on a web page. You need to specify the frame-ancestor attribute, otherwise the browser security policy will prevent the application from rendering in the page.

The frame-ancestor attribute prevents "Clickjacking" attacks by specifying which domains are allowed to provide embedded content to your site. To enable a domain to embed a Cloudera Machine Learning application, set the environmental variable CDSW_FRAME_ANCESTORS to contain one or more websites as follows:

- The name of the embedding website, specified in host-source form.
- You can specify multiple websites as a comma-separated list.

You can set the environmental value in **Project Settings Advanced**.

Resource Profile

1 vCPU / 2 GiB Memory

Environment Variables

CDSW_APP_POLLING_ENDPOINT	/	-	+
CDSW_FRAME_ANCESTORS	http://*.dev.cldr.work, https://*.dev.cldr.w	-	+

Environmental variables will override the [project environment](#).

For more information on host-source form, see: *CSP: frame-ancestors*.

Related Information

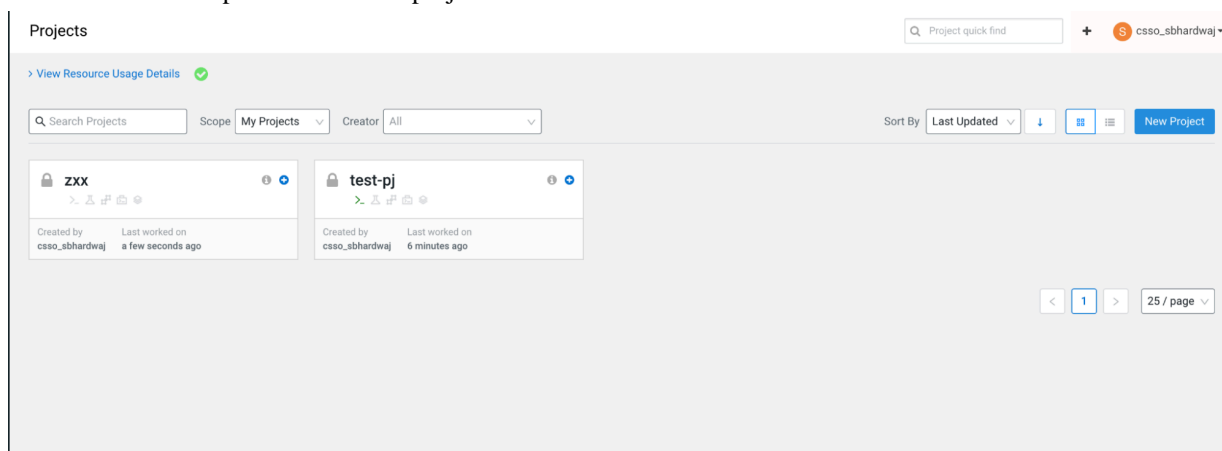
[CSP: frame-ancestors](#)

Setting up Cloudera Machine Learning Workspaces for high volume Workloads

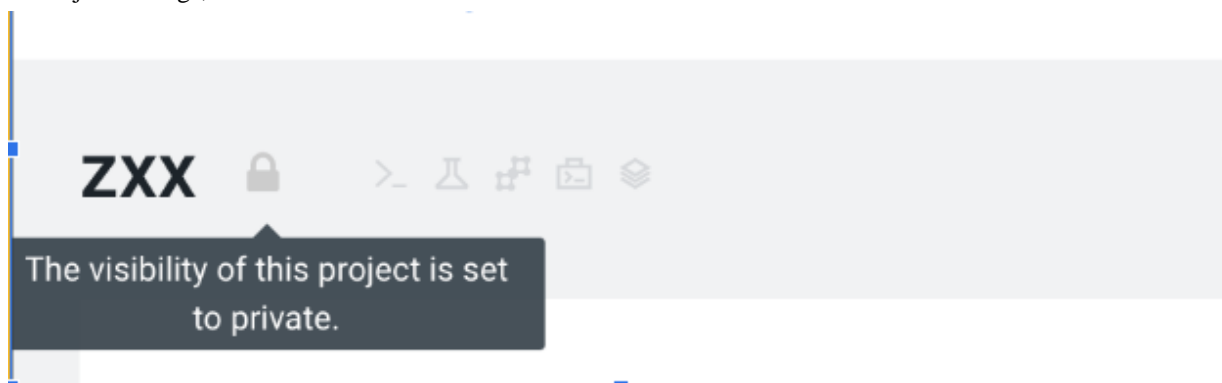
Autoscaling in Cloudera Machine Learning enables seamless scaling up of clusters, accommodating sessions, experiments, model metrics, jobs, and applications with increased user demand. In order to ensure seamless functioning of Cloudera Machine Learning Workspaces with high volume workloads, core Cloudera Machine Learning flows such as workspace suspend, resume, backup, upgrade, and editing are validated with high volume workloads. Additionally, clusters can be efficiently downsized by adjusting the autoscale range in workloads without disrupting control plane capabilities.

To prepare a Cloudera Machine Learning Workspace for high volume workloads, you need to perform the following steps to modify the pod quota limit for the Cloudera Machine Learning tenant, and then do the same for the Cloudera Machine Learning Workspace.

1. Go inside the workspace and create a project.



2. In Project settings, click the lock button.



3. In Advanced settings, add the key `VERRIDE_PODQUOTA` and enter the value for the new pod limit to set.

csso_sbhardwaj / zxx / Project Settings / Advanced

Project Settings

General Runtime **Advanced** SSH Tunnels Data Connections Delete Project

Environment Variables
Set project environment variables that can be accessed from your scripts.
Environment variable **values** are only visible to **collaborators** with **write** or higher access. They are a great way to securely store confidential information such as your AWS or database credentials. Names are available to all users with access to the project.

CDSW_APP_POLLING_ENDPOINT	/	- +
PROJECT_OWNER	csso_sbhardwaj	- +
OVERRIDE_PODQUOTA	1	- +

Submit

Verified flows and configurations

1. Enhanced Pod Count Limit

Scaling up to 250 Pods within a single user namespace is verified, assuming that the workspace possesses sufficient computational resources and storage capacity to accommodate these pods.

2. Auto-Scaling Certification

The Control Plane actions have been verified to scale up to 100 Nodes.

3. Suspension and resumption of 100 Nodes has been successfully tested.

4. Workspace Management

- Backup of workspaces with up to 100 Nodes.
- Upscale and downscale of cluster nodes between 1-100 is now verified by changing the worker node range within the edit section of workspace.
- Successfully tested upgrading a 100 node Workspace from version at 'V-1'.

5. Basic sanity testing of the following workload flows was conducted for workspaces with 100 Nodes:

- Creation of Sessions, Applications, Jobs, Model Metrics and Experiments for a user after workspace creation and upgrading.
- Deletion of Sessions, Applications, Jobs, Model Metrics and Experiments for high volume workloads.
- I/O, memory usage validation during upscale and downscale operations.

Known issues

There are a few known issues that have been identified during the certification process. These issues are currently being addressed.

1. During concurrent creation of sessions via Cloudera Machine Learning Workspace users, some pods may fail to come up, with a ~3% failure rate.
2. For high volume workspaces, Cloudera Machine Learning Workspace applications may fail to restart after modify/resume and upgrade operations. In this case, manually restart the affected application.

Host name required by Learning Hub

Learning Hub requires internet access to link to the displayed content. Learning Hub cannot be supported on a fully airgapped cluster.

The following domain must be added to the allow list so that links from the content will work:

- *.raw.githubusercontent.com