

Machine Learning

ML Workspaces (Private Cloud)

Date published: 2020-07-16

Date modified: 2024-03-05

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

Provision an ML Workspace.....	4
Monitoring ML Workspaces.....	5
Removing ML Workspaces.....	5
Upgrading ML workspaces.....	6
Backing up ML workspaces.....	6
Workspace Backup and Restore Prerequisites.....	7
Back up an ML workspace.....	8
Restore an ML workspace.....	8


Provision an ML Workspace

In CML on Private Cloud, the ML Workspace is where data scientists get their work done. After your Admin has created or given you access to an environment, you can set up a workspace.

Before you begin

The first user to access the ML workspace after it is created must have the EnvironmentAdmin role assigned.

Procedure

1. Log in to the CDP Private Cloud web interface using your corporate credentials or other credentials that you received from your CDP administrator.
 2. Click ML Workspaces.
 3. Click Provision Workspace. The Provision Workspace panel displays.
 4. In Provision Workspace, fill out the following fields.
 - a) Workspace Name - Give the ML workspace a name. For example, test-cml. Do not use capital letters in the workspace name.
 - b) Select Environment - From the dropdown, select the environment where the ML workspace must be provisioned. If you do not have any environments available to you in the dropdown, contact your CDP admin to gain access.
 - c) Namespace - Enter the namespace to use for the ML workspace.
 - d) NFS Server - Select Internal to use an NFS server that is integrated into the Kubernetes cluster. This is the recommended selection at this time.
The path to the internal NFS server is already set in the environment.
 5. In Production Machine Learning, select to enable the following features.
 - a) Enable Governance - Enables advanced lineage and governance features.
Governance Principal Name - If Enable Governance is selected, you can use the default value of mlgov, or enter an alternative name. The alternative name must be present in your environment and be given permissions in Ranger to allow the MLGovernance service deliver events to Atlas.
 - b) Enable Model Metrics - Enables exporting metrics for models to a PostgreSQL database.
 6. In Other Settings, select to enable the following features.
 - a) Enable TLS - Select this to enable https access to the workspace.
 - b) Enable Monitoring - Administrators (users with the EnvironmentAdmin role) can use a Grafana dashboard to monitor resource usage in the provisioned workspace.
 - c) CML Static Subdomain - This is a custom name for the workspace endpoint, and it is also used for the URLs of models, applications, and experiments. Only one workspace with the specific subdomain endpoint name can be running at a time. You can create a wildcard certificate for this endpoint in advance. The workspace name has this format: <static subdomain name>.<application domain>
-  **Note:** The endpoint name can have a maximum of 15 characters, using alphanumeric characters and hyphen or underscore only, and must start and end with an alphanumeric character.
7. Click Provision Workspace. The new workspace provisioning process takes several minutes.

What to do next

After the workspace is provisioned, you can log in by clicking the workspace name on the Machine Learning Workspaces page. The first user to log in must be the administrator.

Related Information

[Monitoring ML Workspaces](#)

[Removing ML Workspaces](#)

Monitoring ML Workspaces

This topic shows you how to monitor resource usage on your ML workspaces.

About this task

Cloudera Machine Learning leverages Prometheus and Grafana to provide a dashboard that allows you to monitor how CPU, memory, storage, and other resources are being consumed by ML workspaces. Prometheus is an internal data source that is auto-populated with resource consumption data for each workspace. Grafana is a monitoring dashboard that allows you to create visualizations for resource consumption data from Prometheus.

Each ML workspace has its own Grafana dashboard.

Before you begin

Required Role: MLAdmin

You need the MLAdmin role to view the Workspace details page.



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

Procedure

1. Log in to the CDP web interface.
2. Click ML Workspaces.
3. For the workspace you want to monitor, click Actions Open Grafana .

Results

CML provides you with several default Grafana dashboards:

- K8s Cluster: Shows cluster health, deployments, and pods
- K8s Containers: Shows pod info, cpu and memory usage
- K8s Node: Shows node cpu and memory usage, disk usage and network conditions
- Models: Shows response times, requests per second, cpu and memory usage for model replicas.

You might choose to add new dashboards or create more panels for other metrics. For more information, see the *Grafana documentation*.

What to do next



Note: Prometheus captures data for the previous 30 days.

Related Information

[Monitoring and Alerts](#)

Removing ML Workspaces

This topic describes how to remove an existing ML workspace and clean up any cloud resources associated with the workspace. Currently, only CDP users with both the MLAdmin role and the EnvironmentAdmin account role can remove workspaces.

Procedure

1. Log in to the CDP web interface.
2. Click ML Workspaces.
3. Click on the Actions icon and select Remove Workspace.
 - a) Force Delete - This property is not required by default. You should first attempt to remove your workspace with this property disabled.

Enabling this property deletes the workspace from CDP but does not guarantee that the underlying kubernetes resources used by the workspace are cleaned up properly. Go to you kuberketes administration console to make sure that the resources have been successfully deleted.
4. Click OK to confirm.

Upgrading CML workspaces and supported upgrade paths

Learn how to upgrade Cloudera Machine Learning (CML) workspaces and the supported upgrade paths.

How to upgrade CML workspaces

1. Log in to CML web interface as an Administrator.
2. On the **Workspace** page, click Upgrade corresponding to the workspace name.

Supported upgrade paths

The following table lists the supported upgrade paths for upgrading CML workspaces:

Current ML version	Target ML version
1.5.2	1.5.3
1.5.1/1.5.1 CHF	1.5.2/1.5.3

Backing up ML workspaces

Cloudera Machine Learning makes it easy to create machine learning projects, jobs, experiments, ML models, and applications in workspaces. The data and metadata of these artifacts are stored in different types of storage systems in private clouds or in external NFS-backed workspaces outside of a private cloud.

You can backup an ML workspace, and restore it later. The backup preserves all files, models, applications and other assets in the workspace (files are not backed up by CML automatically for external NFS-based workspaces). All workspace backups can be viewed in the Workspace Backup Catalog UI.

The Backup and Restore feature gives you the ability to backup all of your data (except files in external NFS-backed workspaces) to protect your machine learning artifacts against disasters. If your Cloudera Machine Learning workspace is backed up, this feature lets you restore the saved data so that you can recover your ML artifacts as they were saved in the desired backup. The Backup and Restore feature gives the administrator the ability to take “on-demand” backups of CML workspaces. Core services running in the workspace are shut down during the backup process to ensure consistency in the backup data. It is recommended that backups are taken during off-peak hours to minimize user impacts.

The time required to complete backing up a workspace depends on the amount of data to copy. The backup process copies data from both block volumes and internal NFS. In general, the time taken to backup NFS is the dominant factor. You should regularly back up CML workspaces.

The time to backup NFS is highly dependent on the amount of data, and on the nature and number of files. Based on the amount of data, you can set a timeout value while taking backup. You can view the status of ongoing/old backups on CML workspace UI and backup catalog UI.

There is currently no restriction on the number of backups one can take, and the backup snapshots are retained indefinitely in the underlying private cloud cluster as long as the original workspace (from which this backup was taken from) is not deleted.. CML workspace backup details are stored in the Workspace Backup Catalog UI in the CML control plane, and these entries may be listed, viewed, deleted or restored as desired.



Note: Deleting workspace backups from UI is not yet supported.

Restoring a backup overwrites the existing CML workspace (from which this backup was taken from) wherein the restored data is automatically imported. All the projects, jobs, applications, etc., that were in existence during the backup are automatically available in the new workspace. Restoring a CML backup overwrites the existing workspace with a new one and then launches restore jobs to create storage volumes from the backup snapshots. The restore process takes longer than a regular workspace provisioning operation due to the extra work in copying data from backup to the new storage volumes. Restores are always full-copy restores. The time to restore is dominated by NFS restoration, which takes at least as long as the time to backup the file system. The restored workspace is always created with the latest CML software version, which may be different from the CML version of the original workspace that was backed up.



Note: ML workspace Backup and Restore feature is available on both ECS and OCP, through the CML UI only.

Workspace Backup and Restore Prerequisites

To backup and restore workspaces, check that the following prerequisites are satisfied.

The following prerequisites apply to Backup functionality.

- Backup is enabled only for workspaces that are successfully installed.
- All workloads (sessions, jobs, applications, models) should be turned off manually by the user before taking backup. This will ensure consistency in the backup data.
- Core services running in the workspace are shut down during the backup process. So, during the backup process, the workspace will not be accessible to the customer.
- It is recommended that backups are taken during off-peak hours to minimize user impacts.
- Time to backup is proportional to the amount of data present in the workspace. So, give sufficient timeout when triggering backup.
- Backup is supported for both external and internal NFS-backed workspaces.
- Make sure enough disk space is available for taking workspace backup. Workspace backup generally takes an equivalent amount of storage space compared to the workspace itself.

The following prerequisites apply to Restore functionality.

- Workspace restore doesn't create a new workspace. It will instead replace the running workspace with an older backup.
- Restore process, overwrites the existing workspace with one of the older backups. This means that anything on the running workspace which is not backed up will get lost forever. So, make sure you really want to restore an older version of the workspace. If you want to save the current state before restore, you can do so by first taking a new backup and then proceeding with the restore.
- All workloads (sessions, jobs, applications, models) should be turned off manually by the user before triggering restore.
- All workloads (sessions, jobs, applications, models) should be turned on manually by the user after restore has completed.
- Time to restore is proportional to the amount of data present in the backup. In general, restoration will take at most 12 hours to complete.

- Always make sure that any ongoing backup for a workspace is completed (by looking at workspace status and backup event logs), before triggering restore for it.

Back up an ML workspace

Backing up an ML workspace preserves all files, models, applications, and other assets in the workspace, although files in external NFS-backed workspaces are not backed up by CML automatically.

Procedure

1. In the Workspaces UI, find the workspace to back up. The workspace must be in the Installation completed state, otherwise backup is disabled.
2. Enter the workspace, and manually stop all workloads (sessions, jobs, applications, and models).
For external NFS backed workspaces, manually back up the configured external NFS data to another location. This manual backup of the NFS data will be used when this particular backup is restored in future. Ignore this step if the workspace is configured with internal NFS, as internal NFS data is backed up and restored automatically by CML.
3. In the Actions menu for that workspace, select Backup Workspace.
4. In the Backup Workspace modal, enter a Backup Name to identify the workspace, and enter an appropriate timeout value.
5. Click Backup to start the process.

Results

The workspace shuts down, and the backup process begins. The workspace state changes to reflect the ongoing backup progress. If necessary, click Cancel to cancel the backup process. The backup process can take some time to complete, depending on the amount of data to copy.



Note: The default timeout is 12 hours. The estimated time to complete a backup (from the cloud provider) is now periodically added to the event logs.

What to do next

Monitoring event logs

You can monitor the progress of the backup process by checking the event logs. In the Actions menu for the workspace, click View Event Logs, and then on the Events & Logs tab, click View Event Logs again for the latest backup event.

When the backup process completes, the workspace enters the Installation completed state again.

If there were issues during backup, appropriate error messages will be displayed in the event logs. However the workspace will recover from failure and will be reverted back to the original state when backup was triggered.

Restore an ML workspace

Restoring a backup overwrites the existing CML workspace (from which the backup was taken from) and automatically imports the restored data. All of the projects, jobs, applications and so on in the original workspace are recreated in the new one.

About this task



Note: Restoring a workspace is a non-reversible operation. The restore process overwrites the existing workspace with older backup data. Any data in the running workspace that is not backed up will be lost. To save the current state, take a new backup before proceeding with the restore operation.

Procedure

1. In the Workspace Backups UI, find the workspace to restore. You can search for the workspace name or CRN. There can be multiple backups for a given workspace.
2. Enter the workspace, and manually stop all workloads (sessions, jobs, applications, and models).
For external NFS backed workspaces, copy the manual backup of external NFS data (corresponding to this particular backup) to the configured external NFS export. Ignore this step if the workspace is configured with internal NFS, as internal NFS data is backed up and restored automatically by CML.
3. Look for the backup to restore, and click Restore. The restore process starts, and the workplace state changes to Creating workspace.

Results

The restore process can take some time, depending on the amount of data to copy. When it is complete, you can find the restored workspace in the Workspaces UI.



Note: If there is an issue during the restore process, the event log will show the relevant error messages. In case of error, the workspace will not recover from the failure automatically and will not revert back to the original state prior to the restore operation.

What to do next

Monitoring event logs

You can monitor the progress of the backup process by checking the event logs. In the Actions menu for the workspace, click View Event Logs, and then on the Events & Logs tab, click View Event Logs again for the latest backup event.

When the backup process completes, the workspace enters the installation completed state again.

If there were issues during backup, appropriate error messages will be displayed in the event logs. However the workspace will recover from failure and will be reverted back to the original state when backup was triggered.