

CML Suspend and Resume Workspace (Preview)

Date published: 2022-10-28

Date modified: 2022-10-28

Legal Notice

© Cloudera Inc. 2022. 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.

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Contents

Legal Notice	2
Contents	3
Suspend and resume ML workspaces	4
CDP Entitlement	4
CDP Cross-Account Role Permissions	4
Suspend and resume a CML workspace	4
Known Issues	5

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Suspend and resume ML workspaces

Cloud consumption costs are a pain point for many public cloud users. The CML Suspend feature allows users to scale down the Kubernetes pods running on CML infra and CPU/GPU nodes for a given ML workspace. When the resume operation is performed on the suspended workspace, the suspended pods scale up.

A suspended CML workspace has all its autoscaling node groups, except the Platform Infra node group, shrunk to zero instances, thereby saving compute instance costs for the duration the workspace is suspended. However, when a workspace is suspended, you cannot access the workspace URL, and all associated models, applications, sessions, and jobs also become unavailable. The suspend operation terminates sessions and jobs, so the suspend operation should be started only after those operations have finished. When the workspace is resumed, models and applications automatically resume operation at the same URLs as before.

A CDP environment might have an automatic garbage collection script that deletes EBS volumes that are not attached to an EC2 instance. EBS volumes are detached when a CML workspace is suspended, and may be inadvertently deleted. This can happen even if the volumes are still claimed by EKS, because the claim is made by tags which might not be checked by the GC script. You should back up a workspace before suspension, and double-check that GC scripts do not delete volumes that are owned by EKS.

CDP Entitlement

The feature is gated behind the `ML_SUSPENDABLE` entitlement. This entitlement must be granted to your CDP account before trying to suspend your workspace.

CDP Cross-Account Role Permissions

AWS: None required.

Azure: None required.

Suspend and resume a CML workspace

1. To suspend a CML workspace, in the Workspaces UI, select **Actions**→**Suspend Workspace** for the workspace to suspend. Then click **OK** to start the suspend process.
2. To resume a CML workspace, in the Workspaces UI, select **Actions**→**Resume Workspace** for the workspace to resume. Then click **OK** to start the resume process.

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Known Issues

1. [DSE-22736](#)

During suspend operation for Azure workspaces, the scale down of node groups does not complete immediately. Rather the actual scale down of the node group happens after the suspend operation has been reported as finished.