

Machine Learning

Release Notes

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

What's New.....	4
Known Issues.....	4
Fixed Issues.....	5

What's New

There are no new Cloudera Machine Learning features for CDP Private Cloud 1.5.3.

[Related Information](#)

[Known Issues and Limitations](#)

Known Issues

You might run into some known issues while using Cloudera Machine Learning on Private Cloud.

DSE-35013: First CML workspace creation fails

On RHEL 8.8, during the first CML workspace installation on GPU with Embedded Container Service external registry, pods might get stuck in the init or crashloop state.

First-time workspace installation is expected to fail. Cloudera recommends that you consider this as a test workspace, and apply the following manual workaround for creating subsequent workspaces:

1. Restart or delete the pods which are in init or crashloop state in the test workspace.
2. Once all pods are in the running state, create new workspaces as needed.
3. Delete the test workspace from the CML UI if no longer needed.

DSE-32887: Runtime Updates can not be enabled from the Site administrator UI on Private Cloud Data Services

With this bug present on Private Cloud Data Services 1.5.2 and 1.5.3, Runtime repositories added on the Site Administration>Runtimes page will not be scanned for new Runtimes. Adding Runtime repositories to CML was a new feature in 1.5.2. Unfortunately with this bug, the feature is completely unusable.

Workaround:

Manually enable Runtime updates with an API call using the legacy API key of the administrator user. That will let CML check the registered runtime registries for new runtimes every 24 hours. The syntax of the API call is as follows:

```
curl 'https://<CML_URL>/api/v1/site/config' \  
-X 'PATCH' \  
-u <ADMIN_APIv1_KEY>: \  
-H 'content-type: application/json' \  
-d '{"enable_runtime_updates":true}'
```

DSE-32730: Workloads unable to start up after changing default hadoopCLI addon

Changing the default Hadoop CLI Runtime Addon causes jobs, models, and application workloads to be unable to start up.

Workaround:

1. Open affected workload settings.
2. Update the workload (this updates the Hadoop CLI Addon associated with the workload to the default one.)
3. For Jobs: update.
4. For Applications: update and restart.
5. For Models: deploy a new build.

Please see Machine Learning Site administration>Disable Addons for related tasks.

DSE-28768: Spark Pushdown is not working with Scala 2.11 runtime

Scala and R are not supported for Spark Pushdown.

Workaround: None.

Related Information

[Known Issues and Limitations](#)

Fixed Issues

This section lists issues fixed in this release of Cloudera Machine Learning (CML) on Private Cloud.

DSE-32639: Installation and upgrades failing due to environment name character restriction

In Private Cloud Data Services 1.5.2, the supported length for environment names was restricted between 5 and 28 characters for creating and upgrading workspaces, causing failures. This issue has been fixed. You no longer need to restrict the environment name length to 28 characters.

DSE-34100: Migration readiness check fails on BuildKit-enabled cluster

Earlier, the migration readiness checks failed on clusters having BuildKit instead of Docker. This issue has been fixed.

DSE-31914: The model metrics ingress configuration does not support TLS

This issue has been fixed.

DSE-32250: Spark executor usage is not reported correctly

Earlier, when you exported the Spark executor usage data from the `Site Admin Usage` tab, the exported data in the report was same as the session in which the executor was launched. This included incorrect CPU, memory, create time, and stop time details of the Spark executors. This issue has been fixed.

DSE-33098: Unable to synchronize teams in CML without having an MLAdmin user in that team

Earlier, CML required that each team to have an MLAdmin user to manage the team. This is no longer a hard requirement. You can create teams of users with MLUser-only roles.

DSE-33765: CML workspace upgrades failing due to buggy logic in fetching the Helm values

This issue has been fixed.

DSE-34088: Unable to view Spark UI executor logs in CML

This issue has been fixed.