

Data Lake Scaling (Preview)

Date published: 2022-05-16

Date modified: 2022-05-17

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.

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
Overview	4
Limitations	4
Pre-requisites	5
Scaling the Data Lake through the CDP UI	7
Post-requisites	8
Scaling the Data Lake through the CDP CLI	8
Post-requisites	9
Recovering after a failed scaling operation	10

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.

Overview

Data Lake scaling is the process of scaling up a light duty Data Lake to the [medium duty form factor](#), which has greater resiliency than light duty and can service a larger number of clients. You can trigger the scale-up in the CDP UI or through the CDP CLI.

During a typical Data Lake scaling operation, the metadata maintained in the Data Lake services is automatically backed up, a new medium duty Data Lake is created within the environment, and the Data Lake metadata is automatically restored to the new medium duty Data Lake.

Note: RAZ-enabled Data Lakes are not currently eligible for automatic restore during a scaling operation; they will be automatically backed up, but must be manually restored after the scaling is complete. If RAZ is in use, before starting the Data Lake backup make sure that the appropriate Ranger policy exists with access to the backup location in the cloud. See instructions for manually restoring a RAZ-enabled Data Lake [here](#).

Supportability Matrix

	CDW	CDE	CML	ODX	CDF
NON-RAZ	Yes	Yes	Yes	Yes	No
RAZ	No	Yes	Yes	Yes	No

Limitations

1. Scaling is only supported for CDP Runtime versions 7.2.7 and above, because medium duty Data Lakes are not supported for earlier versions.
2. Only resizing from a light duty Data Lake to a medium duty HA Data Lake is supported.
3. Data Lake scaling is not supported for GCP environments.
4. The scaling operation requires an outage and should be performed during a maintenance window. No metadata changes may occur during the scaling, as these changes will no longer be present once the scaling operation completes

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.

(the previously backed up metadata is being restored). Suspend any operations that may result in any SDX metadata change during the scaling operation.

5. Data Hub clusters should be stopped before the scaling operation begins. For any cluster that cannot be stopped, stop all of the services on the Data Hub through the Cloudera Manager UI.
6. CDW jobs using the warehouses will fail during the scaling operation.
7. CDF flows should be re-created after a scaling operation.
8. Currently, FreeIPA scaling is not supported.

Pre-requisites

Prior to scaling the Data Lake, ensure that the following are in place:

1. The Data Lake must be running to perform the scaling operation.
2. For RAZ-enabled Data Lakes, update the appropriate Ranger policy to give the backup and restore feature permission to access the backup location in the cloud. See instructions for configuring RAZ for backup [here](#).
3. Make sure that Atlas is up to date and has processed all the lineage data in Kafka. To do this, please follow the steps below:
 - a. **ssh** into the master node of your light duty Data Lake
 - b. Switch to the super user for the node by running **sudo su**
 - c. Copy over the following script into a file called **check_atlas_updated.sh**

```
#!/usr/bin/env bash

# Determine Atlas keytab path.
ATLAS_KT=$(find / -wholename "*atlas-ATLAS_SERVER/atlas.keytab"
2>/dev/null | head -n 1)

# Setup required configuration files if needed.
if [[ ! -f jaas.conf ]]; then
    ATLAS_PRINCIPAL=$(klist -kt "${ATLAS_KT}" | grep -o -m 1
"atlas\AS*")
    printf "KafkaClient {
\tcom.sun.security.auth.module.Krb5LoginModule required
\tuseKeyTab=true
\tkeyTab=\"%s\"
\tprincipal=\"%s\";\n};\n" "${ATLAS_KT}" "${ATLAS_PRINCIPAL}"
> jaas.conf
fi

if [[ ! -f client.config ]]; then
```

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.

```

        printf
"security.protocol=SASL_SSL\nsasl.kerberos.service.name=kafka\n" >
client.config
fi

# Determine the Kafka bootstrap server to use.
KAFKA_SERVER=$(grep --line-buffered -oP
"atlas.kafka.bootstrap.servers=\K.*" \
    /etc/atlas/conf/atlas-application.properties | awk -F',' '{print $1}')

# Export Kafka-specific environment variables.
export KAFKA_HEAP_OPTS="-Xms512m -Xmx1g"
export
KAFKA_OPTS="-Djava.security.auth.login.config=${PWD}/jaas.conf"

# Kinit into Atlas keytab as Atlas user.
kinit -kt "$ATLAS_KT" "atlas/$(hostname -f)" 2>/dev/null

# Obtain Atlas lineage information.
LINEAGE_INFO=$(/opt/cloudera/parcels/CDH/lib/kafka/bin/kafka-consumer-groups.sh \
    --bootstrap-server "${KAFKA_SERVER}" --describe --group atlas \
    \
    --command-config="${PWD}/client.config" 2>/dev/null \
    | awk '{print $2, $6}')

if [[ -z "$LINEAGE_INFO" ]]; then
    echo "*ERROR*: Unable to get lineage info for Atlas. Please look
at the created configuration files to make sure they look correct."
    exit 1
fi

# Parse lineage information and determine if Atlas is out of date.
LINEAGE_LAG_VALS=($LINEAGE_INFO)
NUM_LAG_VALS=${#LINEAGE_LAG_VALS[@]}
OUT_OF_DATE_TOPICS=""
for (( i = 2; i < ${NUM_LAG_VALS}; i += 2 )); do
    if [[ ${LINEAGE_LAG_VALS[$i] + 1} != '-' &&
${LINEAGE_LAG_VALS[$i] + 1} != '0' ]]; then

OUT_OF_DATE_TOPICS="${OUT_OF_DATE_TOPICS}${LINEAGE_LAG_VALS[$i]}, "
    fi
done

if [[ -z "$OUT_OF_DATE_TOPICS" ]]; then

```

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.

```

        echo "Atlas is up to date! Feel free to continue with the
migration."
else
        echo "The following Atlas topics are not up to date:
${OUT_OF_DATE_TOPICS%??}!"
        echo "Please wait until Atlas is entirely up to date before
continuing with the migration."
fi
    
```

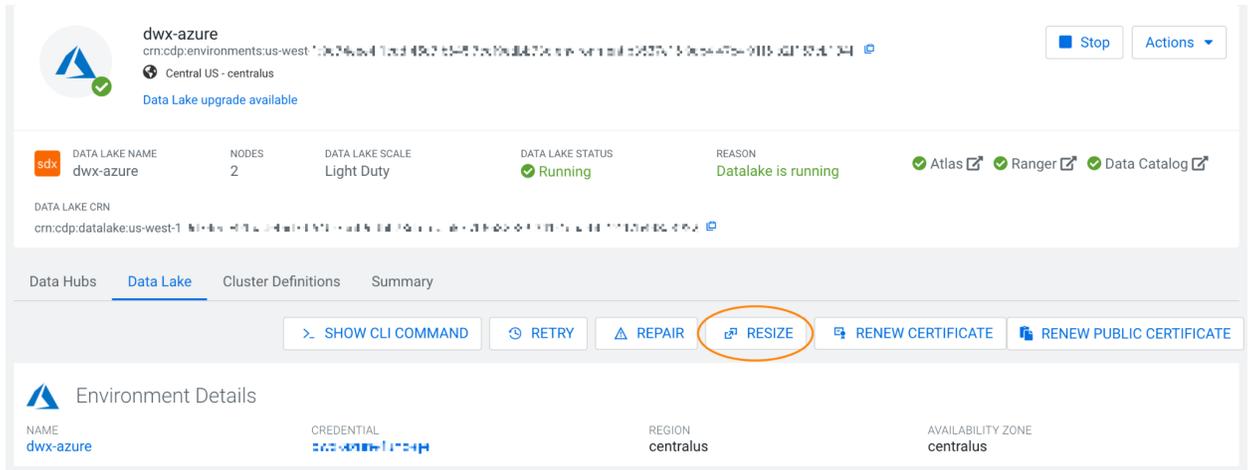
- d. Allow the new script to be run by running **chmod +x check_atlas_updated.sh**
 - e. Run the script with **./check_atlas_updated.sh**. The script will tell you if Atlas is up to date or not. If it isn't, wait a while and check again. You should only begin the scaling process if the script tells you that Atlas is up to date.
4. If you are using CDW, stop the virtual warehouses and data catalogs associated with the environment.
 5. If you are using CDE, do one of the following:
 - a. Upgrade to CDE 1.15.
 - b. Create new service.
 - i. Take a backup of your jobs following [Backing up Cloudera Data Engineering jobs](#).
 - ii. Create a new DE service and virtual cluster.
 - iii. Restore the jobs following the instructions in [Restoring Cloudera Data Engineering jobs from backup](#).

Scaling the Data Lake through the CDP UI

Required role: *EnvironmentAdmin* or *Owner* of the environment

1. Stop all of the attached Data Hub clusters that can be stopped, to make sure that there are no changes to HMS metadata during the scaling operation. For any cluster that cannot be stopped, stop all of the services on the Data Hub through the Cloudera Manager UI.
2. Verify that the `DATALAKE_ADMIN_ROLE`, `RANGER_AUDIT_ROLE`, and `LOG_ROLE` have read/write permissions to the backup location. See the [Data Lake backup and restore documentation](#) for more information on these permissions. `LOG_ROLE` is specific to [Data Lake restore](#).
3. In the CDP UI, click **Data Lakes** and select the Data Lake that you want to scale.
4. Click **Resize**.

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.



You will be asked to confirm that you want to resize the Data Lake, after which the scaling process will begin. The scaling operation is finished when the Data Hub clusters have been automatically refreshed, which happens after the original light duty Data Lake has been deleted. Check the Event History to verify that the Data Hubs have been refreshed.

5. When the scaling operation is finished, if you are upscaling a RAZ-enabled Data Lake, [manually restore the Data Lake](#) from the backup taken by the scaling process. To restore, you will need to add the `restore_to_raz` policy in Ranger and then run the `cdp_datalake_restore-datalake` command in the CDP CLI.

Post-requisites

After the scaling process is complete, complete the following tasks:

1. If RAZ is not being used, resync the ID broker mappings to the Data Lake.
2. Start the data catalogs and virtual warehouses.
3. Restart or re-deploy any existing CML workloads:

CML Entity	Action Needed
Model	Re-deploy
Applications	Re-deploy
Experiments	Re-deploy
Sessions	Restart New
Jobs	Restart
Notebook	Re-deploy

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.

4. Re-create any CDF flows.
5. Start the Data Hub cluster services, if you stopped them before the scaling operation.

Data Hubs that were stopped before the scaling operation should continue to work when the scaling completes, by communicating with the new Data Lake automatically when they are re-started.

Scaling the Data Lake through the CDP CLI

Required role: *EnvironmentAdmin* or *Owner* of the environment

1. Stop all of the attached Data Hub clusters that can be stopped, to make sure that there are no changes to HMS metadata during the scaling operation. For any cluster that cannot be stopped, stop all of the services on the Data Hub through the Cloudera Manager UI.
2. Verify that the DATALAKE_ADMIN_ROLE, RANGER_AUDIT_ROLE, and LOG_ROLE have read/write permissions to the backup location. See the [Data Lake backup and restore documentation](#) for more information on these permissions. LOG_ROLE is specific to [Data Lake restore](#).
3. To trigger scaling from the CDP CLI, run the `cdp datalake resize-datalake` command. For example:

```
cdp datalake resize-datalake --datalake-name <mydatalake>
--target-size MEDIUM_DUTY_HA
```

Option	Description
<code>--datalake-name</code>	Name or CRN of the Data Lake that you want to upscale.
<code>--target-size</code>	Currently only MEDIUM_DUTY_HA is supported.

4. Monitor the Event History. The scaling operation is finished when the Data Hub clusters have been automatically refreshed, which happens after the original light duty Data Lake has been deleted. Check the Event History to verify that the Data Hubs have been refreshed.
5. When the scaling operation is finished, if you are upscaling to a RAZ-enabled Data Lake, [manually restore the Data Lake](#) from the backup taken by the scaling

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.

process. To restore, you will need to add the `restore_to_raz` policy in Ranger and then run the [`cdp_datalake_restore-datalake`](#) command in the CDP CLI.

Post-requisites

After the scaling process is complete, complete the following tasks:

1. If RAZ is not being used, resync the ID broker mappings to the Data Lake.
2. Start the data catalogs and virtual warehouses.
3. Restart or re-deploy any existing CML workloads (Sessions, Jobs, Applications, Experiments, notebooks, and Models).
4. Re-create any CDF flows.
5. Start the Data Hub cluster services, if you stopped them before the scaling operation.

Data Hubs that were stopped before the scaling operation should continue to work when the scaling completes, by communicating with the new Data Lake automatically when they are re-started.

Recovering after a failed scaling operation

You can recover from a failed scaling operation by returning a Data Lake to its original state before the resize operation was started. Be sure that this is what you want to do before proceeding.

Data Lake recovery simply reattaches and starts the original light duty Data Lake in the environment. All of the instances, disks, and databases are unchanged from their original state.

Currently the recovery logic uses logs that are cleaned up every 24 hours, so recovery may not be possible after this time period. Try to recover as soon as an error is noticed.

If recovery cannot be started, or fails for any reason, reach out to the Cloudera support team, who can manually recover your Data Lake.

1. Trigger the recovery command through the CDP CLI:

```
cdp_datalake_recover-datalake --datalake-name <mydatalake>
```

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.

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.