CDP Private Cloud Data Services Data Recovery

Date published: 2023-12-16 Date modified: 2024-03-05



Legal Notice

© Cloudera Inc. 2025. 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

Backup and Restore Manager in CDP Private Cloud Data Services Management Console		
Learn how backup and restore events work in Backup and Restore		
Manager	4	
Access Backup and Restore Manager in Management Console	5	
Backup Overview section on Management Console Dashboard		
View Backup and Restore Manager	6	
Backups tab in Backup and Restore Manager	6	
Restores tab in Backup and Restore Manager	8	
Creating backup of Control Plane and restoring it	10	
DRS automatic backups (technical preview)	20	
Configuring external storage in ECS for DRS automatic backups		
Initiating DRS automatic backups		
CLI reference for using DRS on Control Plane	24	
Troubleshooting Backup and Restore Manager	25	
CDP Control Plane UI or the Backup and Restore Manager becomes inaccessible after a failed restore event?		
Timeout error appears in Backup and Restore Manager		
Timeout error during backup of OCP clusters		
Stale configurations in Cloudera Manager after a restore event		
Restore event for an environment backup fails with an exception		
Using DRS with CDW	31	

Data Recovery Service overview

The Data Recovery Service (DRS) is a microservice in CDP Private Cloud Data Services that allows you to back up and restore Kubernetes namespaces and resources on both Embedded Container Service (ECS) and OpenShift Container Platform (OCP). You can backup and restore namespaces for Control Plane and other supported services. The "Backup and Restore Manager" in the CDP Private Cloud Data Services Management Console, and the CDP CLI options for the respective services leverage the DRS capabilities to run the backup and restore events (jobs).

Cloudera recommends that you create a backup of your Kubernetes namespace before a maintenance activity, before you upgrade, or in general, as a best practice.

Role required: PowerUser

By default, DRS is located in the [***CDP_INSTALLATION_NAMESPACE***]-drs namespace. For example, if the CDP Private Cloud Data Services installation is located in the *cdp* namespace, the *drs* namespace is automatically named *cdp-drs*. If you have multiple CDP Private Cloud Data Services installations (as in OCP), DRS is named accordingly.

When you initiate the backup event in the Backup and Restore Manager for Control Plane, DRS takes a backup of the following resources and data:

- Kubernetes resources associated with the cdp namespace and the embedded vault namespaces of the Control Plane
 in CDP Private Cloud Data Services. The resources include deployment-related information, stateful sets, secrets,
 and configmaps.
- Data used by the stateful pods, such as the data in the embedded database and Kubernetes persistent volume claim.

DRS requires CSI snapshots to back up and restore Kubernetes namespaces and resources. The CSI snapshots are enabled on ECS by default. You might require an additional license to enable CSI snapshots in Red Hat ODF storage on OCP.

Alternatively, you can use the CDP CLI options in CDP CLI drscp.

Learn how backup and restore events work in DRS

The "Backup and Restore Manager "in the CDP Private Cloud Data Services Management Console and the CDP CLI options use Data Recovery Service (DRS) to backup and restore Kubernetes namespaces and resources on Embedded Container Service (ECS) and OpenShift Container Platform (OCP) for Control Plane and other supported services. You can take multiple backups and restore any backup as necessary.

Backup event

The backup event does not have any downtime impact and you can backup the Control Plane while it is running.

When you create a backup, DRS:

- 1. initiates the backup event or job for the chosen backup entity. For example, the Control Plane in CDP Private Cloud Data Services,
- 2. assigns an ID called backupCrn to the backup event,
 - The backupCRN appears in the CRN column on the Backup and Restore Manager Backups tab. Click the **CRN** to view more details about the backup event on the Backup [***NAME OF BACKUP***] modal window.
- **3.** creates a backup of the persistent volume claim (PVC) snapshots of the Control Plane namespaces and the backup event's PVC.



Tip: DRS is a microservice in CDP Private Cloud Data Services that backs up and restores the Kubernetes namespaces and resources of supported services. Backup and Restore Manager leverages the DRS capabilities to backup and restore namespaces in Management Console.

Restore event



Important: Do not delete the [***CDP_INSTALLATION_NAMESPACE***]-drs namespace while the restore event is in progress.

When you start the restore event, DRS:

- 1. initiates the restore event based on the chosen backup,
- 2. assigns an ID called restoreCrn to the restore event,

The restoreCRN appears as CRN on the Backup and Restore Manager Restores tab. Click the **CRN** to view more details about the restore event.

3. deletes the existing resources and data,

During this stage of the restore event, the ECS restore vault is sealed and the POD is down which might appear as a failure in the control plane environment. After the restore event is complete, the vault and POD are autorecovered and restored. Depending on the number of resources and data, this step might take a maximum of 10 minutes to complete.

4. restores the resources and data from the backup.

The restore event has a downtime impact because the pods and data are recreated.

Related Information

CDP CLI drscp

Access Backup and Restore Manager in Management Console

To backup and restore Kubernetes namespaces and resources on Embedded Container Service (ECS) and OpenShift Container Platform (OCP), you can use the Backup and Restore Manager in the CDP Private Cloud Data Services Management Console.

You can view the list of available backup entities in the "Backup Overview" section on the Dashboard page of the CDP Private Cloud Data Services Management Console. When you click "View Details" in this section, the "Backup and Restore Manager" appears. You can create backups, and view and perform actions on the available backup and restore entities in the "Backups" tab and "Restores" tab respectively.



Tip: Data Recovery Service (DRS) is a microservice in CDP Private Cloud Data Services that backs up and restores the Kubernetes namespaces and resources of supported services. Backup and Restore Manager leverages the DRS capabilities to backup and restore namespaces in Management Console.

Backup Overview section on Management Console Dashboard

You can view all the available backup entities for all the supported services, total number of backups available for each entity, and the status of each backup event in the "Backup Overview" section on the "Dashboard" page in CDP Private Cloud Data Services Management Console.

The following columns appear in the Dashboard Backup Overview section:

Column name	Description
Backup Entity	Lists all the backup entities that are available for backup. For example, the CDP Control Plane: [***NAMESPACE***] is the Control Plane backup entity.
Total Backups	Total number of backups for the backup entity.

Column name	Description
	Each dot pertains to the current status of a backup event or job. Hover over a dot to view the backup event start timestamp and its current status.
	The yellow dot denotes NOT_STARTED or IN_PROGRESS backup event state, the green dot denotes COMPLETED event state, and the red dot denotes PARTIALLY_FAILED or FAILED event state.

When you click View Details, the **Backup and Restore Manager** appears.

The following sample image shows the **Backup Overview** section:



View Backup and Restore Manager

The Backup and Restore Manager appears after you click "View Details" in the "Backup Overview" section on the "Dashboard" page in CDP Private Cloud Data Services Management Console. The "Backups" tab lists all the backups and the "Restores" tab lists all the restore events.

Click New Backup on the **Backup and Restore Manager** page to initiate a backup event. The page shows the backup entity name and provides the following tabs:

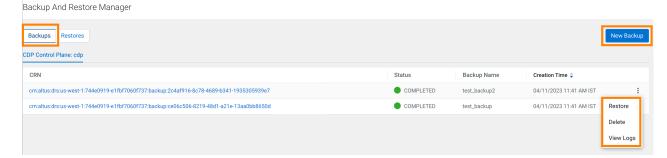
Backups tab in Backup and Restore Manager

The "Backups" tab in Backup and Restore Manager on the CDP Private Cloud Data Services Management Console lists all the available backups. You can create backups, or perform actions such as restore, delete, or view logs for each back up as necessary.

The following table lists the columns that appear on the **Backups** tab:

Column name	Description
CRN	Automatically assigned ID or backupCrn for the backup event. Customer Resource Number (CRN) is the <i>Cloudera</i> -specific identifier provided for the event/job.
	Click the CRN to view more details about the event on the Backup [***NAME OF BACKUP***] modal window.
Status	Current backup event status.
	The event states include NOT_STARTED, IN_PROGRESS, COMPLETED, PARTIALLY_FAILED, and FAILED.
Backup Name	Unique name given to the backup event while initiating the backup event.
Creation Time	Timestamp when the backup event was initiated.

The following sample image shows the **Backups** tab on the **Backup and Restore Manager** page:



You can perform the following actions on each successful backup event:

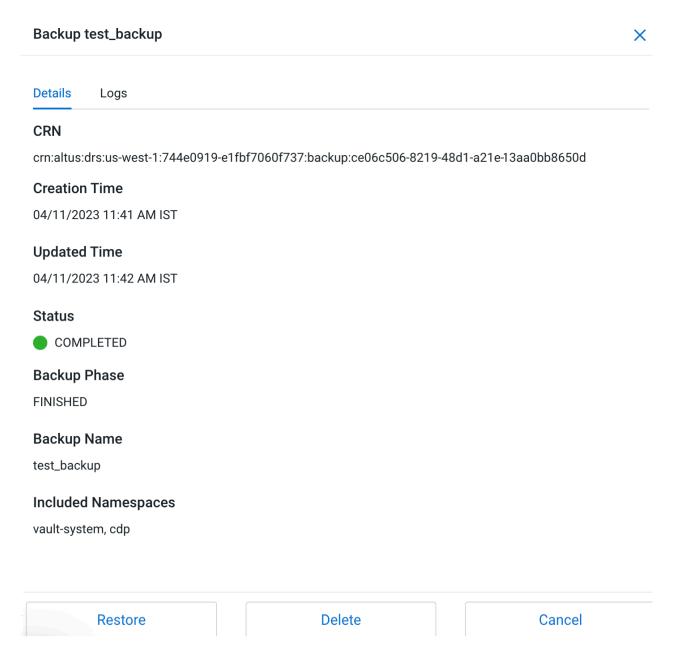
- Restore the backup.
- Delete the backup. This deletes the backup permanently.
- View Logs opens the Backup [***NAME OF BACKUP***] modal window.

Backup [***NAME OF BACKUP***] modal window

On the Backup [***NAME OF BACKUP***] modal window, you can choose to Restore the backup, Delete the backup, or click Cancel to close the window. The window also shows the following tabs:

Tab	Description
Details	 CRN of the backup event Creation Time and date Completed or Updated Time and date Current Status of the backup event Backup Phase the event is running in (such as in-progress or finished) Backup Name that was assigned to the backup event during creation The Included Namespaces in the backup event
Logs	Provides the log details about the backup event.

The following sample image shows the Backup [***NAME OF BACKUP***] modal window:



Restores tab in Backup and Restore Manager

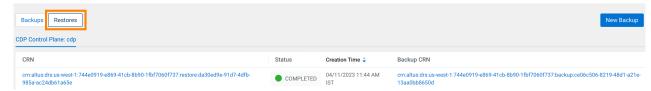
The "Restores" tab in Backup and Restore Manager on the CDP Private Cloud Data Services Management Console lists all the available restore events.

The following table lists the columns that appear on the **Restores** tab:

Column name	Description
CRN	Automatically assigned ID or restoreCrn for the restore event. When you click the CRN, the Restore Details modal window appears.
Status	Current restore event status as COMPLETED or FAILED.
Creation Time	Timestamp when the restore event is initiated.
Backup CRN	CRN of the backup event that is being restored. When you click the Backup CRN , the Backup [***NAME OF BACKUP***] modal window appears.

The following sample image shows the **Restores** tab on the **Backup and Restore Manager** page:

Backup And Restore Manager



Restore Details modal window

When you click the CRN of a restore event, the following tabs appear on the Restore Details modal window:

Column name	Description
Details	 CRN of the restore event Creation Time and date of the restore event Completed or Updated Time and date of the restore event Current Status of the restore event The Restore Phase the event is running in (such as in-progress, pending, failed, or finished) Associated Backup CRN of the backup event that was restored The Included Namespaces in the restore event Warnings or Errors. When a warning appears, you can continue to use the backup or restore event. However, it is advisable to scrutinize the warning to avoid any potential issues. Errors appear if the restore event has failed.
Logs	Provides the log details about the backup event.

The following sample image shows the Restore Details modal window:

Restore Details



Details

Logs

CRN

crn:altus:drs:us-west-1:017a9c10-d8e7-472f-88ed-158946a2fe84:restore:a6668773-defb-4f45-b65dc2160055205a

Creation Time

04/27/2023 5:40 PM IST

Updated Time

04/27/2023 5:45 PM IST

Status



COMPLETED

Restore Phase

FINISHED

Associated Backup CRN

crn:altus:drs:us-west-1:017a9c10-d8e7-472f-88ed-158946a2fe84:backup:a9bde036-5071-4f8f-afd7-141edffb7f9f

Included Namespaces

vault-system, cdp

Creating backup of Control Plane and restoring it

The Backup and Restore Manager in the CDP Private Cloud Data Services Management Console helps you to backup and restore Kubernetes namespaces and resources on Embedded Container Service (ECS) and OpenShift Container Platform (OCP). You can also restore and delete the backups.

Before you begin

Ensure that the following prerequisites are complete:

- You must have the *PowerUser* role.
- For OCP, ensure that a VolumeSnapshotClass is installed with a CSI driver that matches the CSI driver for the storage class used.



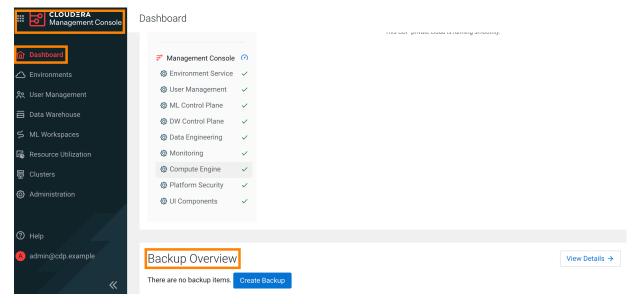
Caution: Multiple restores consume disk space. Therefore, Cloudera recommends that you identify and remove the (restore) replicas, that are not required, manually on a regular basis.

About this task

The following steps show how to create a backup of the Kubernetes namespaces and resources in the Control Plane, restore a backup, delete a backup, view logs for an event, and sample CDP CLI commands.

Procedure

1. Go to the CDP Private Cloud Data Services Management Console Dashboard Backup Overview section.



2. To create a backup, perform the following steps:

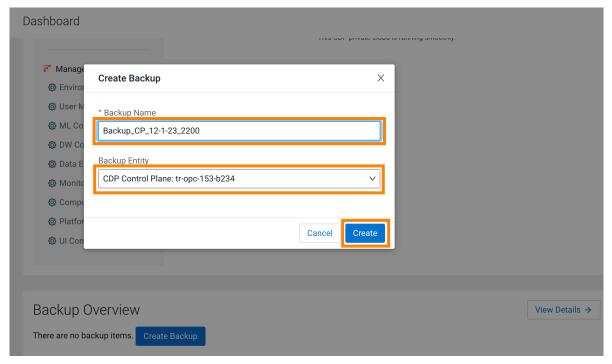
When you create a backup of the Control Plane, Data Recovery Service (DRS) initiates the backup event or job for the chosen backup entity, assigns an ID called backupCrn to the backup event, and creates a backup of the

persistent volume claim (PVC) snapshots of the Control Plane namespaces and the backup event's PVC. CRN or Customer Resource Number is the Cloudera-specific identifier provided for an event or job.

a) Click Create Backup in the Backup Overview section to create the first backup.



b) Enter a unique Backup Name and choose the Backup Entity that you want to back up in the **Create Backup** modal window, and then click Create.



c) DRS initiates the backup event and generates a backupCRN which is an automatically assigned ID for the backup event.

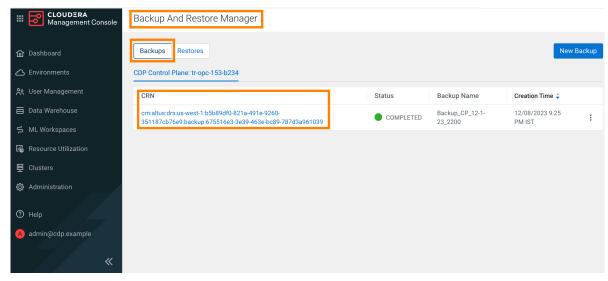
The backup event appears in the **Backup Overview** section.



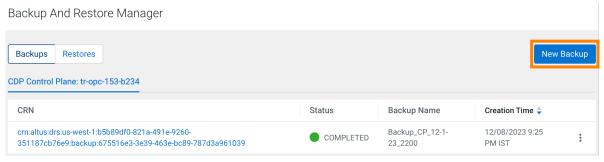
d) Click View Details. The **Backup and Restore Manager** page appears.



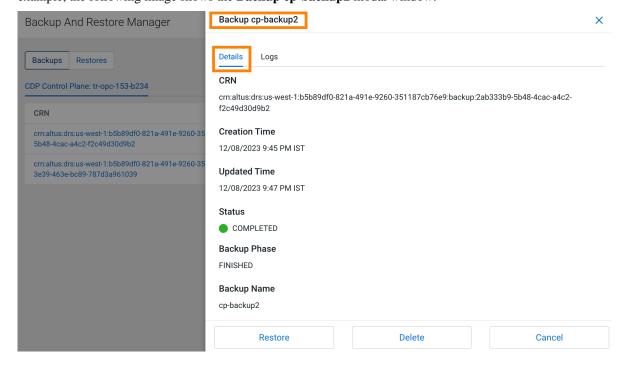
e) The backupCRN appears as a CRN on the **Backup and Restore Manager** page that you can click to view the backup event details.



f) For subsequent backups, click New Backup on the Backup and Restore Manager page.



g) Click the CRN to view more details about it on the Backup [***BACKUP NAME***] modal window. For example, the following image shows the **Backup cp-backup2** modal window.



3. To restore a backup, perform the following steps:

When you start the restore a backup, DRS initiates the restore event based on the chosen backup, assigns an ID called restoreCrn to the restore event, deletes the existing resources and data, and restores the resources and data from the backup.



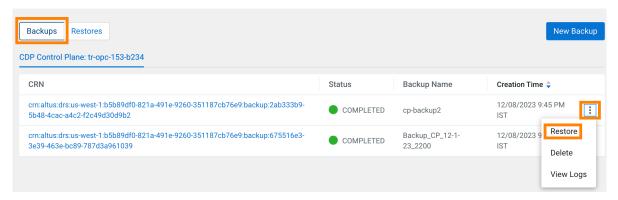
Note: The restore event has a downtime impact because the PODs and data are recreated. During the restore event, the ECS restore vault is sealed and the POD is down which might appear as a failure in the control plane environment. After the restore event is complete, the vault and POD are auto-recovered and restored. Depending on the number of resources and data, this step might take a maximum of 10 minutes to complete. If the environment does not come up, see the logs to troubleshoot. You can also contact your Cloudera account team.

- a) Go to the Backup and Restore Manager Backups tab.
- b) Click Actions Restore, and then click OK in the **Restore** modal window to acknowledge that you want to restore the backup.



Important: Do not delete the [****CDP_INSTALLATION_NAMESPACE***]-drs namespace while the restore event is in progress. For example, if the CDP Private Cloud Data Services installation is located in the *cdp* namespace, the *drs* namespace is automatically named *cdp-drs*.

Backup And Restore Manager





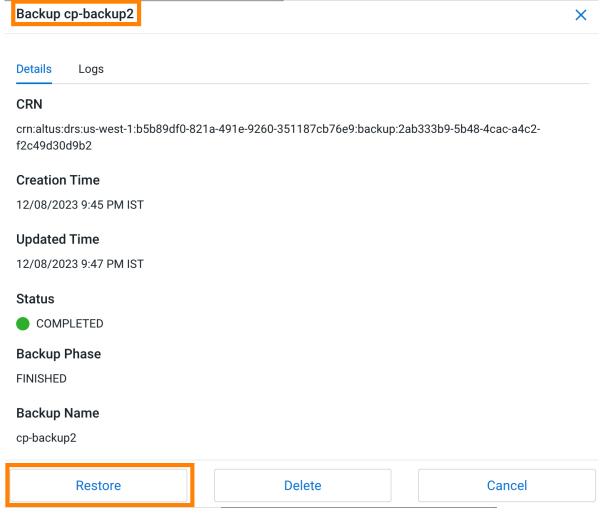
Restore

Are you sure you want to restore this record?

Note: Restore operation will take some time and cause Management UI downtime.

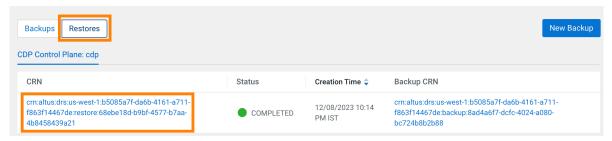


c) Alternatively, click the CRN of the required backup, click Restore on the Backup [***NAME OF BACKUP***] modal window, and then click OK to acknowledge that you want to restore the backup.

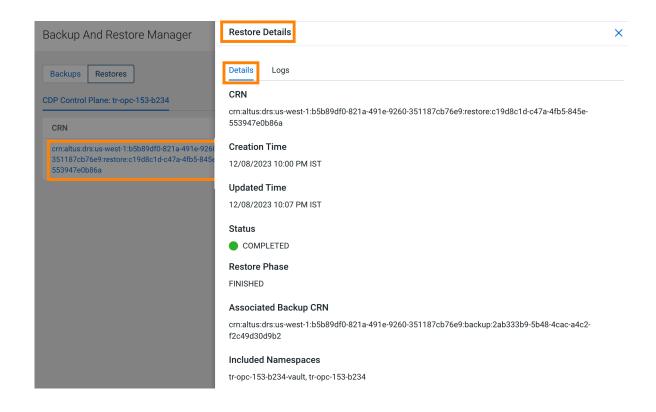


d) Go to the Restores tab to view the CRN for the restore event and other details about the restore event.

Backup And Restore Manager

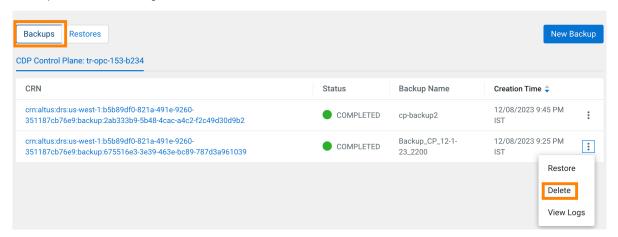


e) Click the CRN for a restore event to see its details on the **Restore Details** modal window.



- **4.** To delete a backup, perform the following steps:
 - a) Go to the Backup and Restore Manager Backups tab.
 - b) Click Actions Delete, and then click OK in the Delete modal window to acknowledge that you want to delete the backup.

Backup And Restore Manager



(!)

Delete

Are you sure you want to delete this record?

Note: You cannot undo this action once performed.

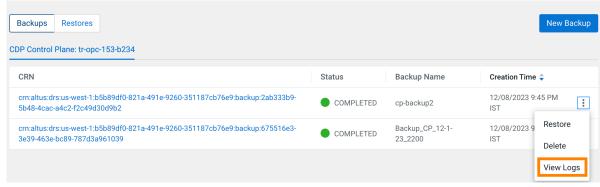


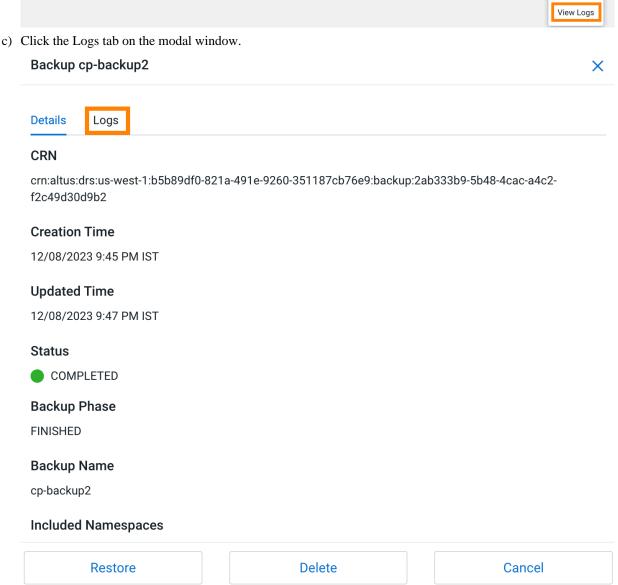
c) Alternatively, click the CRN of the required backup. Click Delete on the Backup [***NAME OF BACKUP***] modal window, and then click OK to acknowledge that you want to delete the backup.

Backup cp-backup2		×	
Details Logs			
CRN			
crn:altus:drs:us-west-1:b5b89df0-821a- f2c49d30d9b2	491e-9260-351187cb76e9:backup::/	2ab333b9-5b48-4cac-a4c2-	
Creation Time			
12/08/2023 9:45 PM IST			
Updated Time			
12/08/2023 9:47 PM IST			
Status			
COMPLETED			
Backup Phase			
FINISHED			
Backup Name			
cp-backup2			
Included Namespaces			
Restore	Delete	Cancel	

- **5.** To view the logs for a backup or restore event, perform the following steps:
 - a) Go to the Backup and Restore Manager Backups tab.
 - b) Click Actions Logs for the required backup.

Backup And Restore Manager





d) Alternatively, you can click the CRN for a backup event on the **Backups** tab, or click the CRN for a restore event on the Restores tab to open the modal window to view the logs for the event.

- **6.** The following sample CDP CLI options show how to create a backup, restore or delete it, and monitor the progress of the events:
 - a) Create a backup using the create-backup CDP CLI option.

The following sample snippet creates a backup named Backup 2.

```
cdp.sh --form-factor private --endpoint-url https://cle-cp1.apps.srdd-os
-01.kcloud.cloudera.com drscp create-backup --backup-name "Backup 2"
```

b) Track the progress of the current status of the specified backupCrn (backup event) using the describe-backup CDP CLI option.

The following sample snippet output shows the current status of the *crn:altus:drs:us-west-1:18be-4c75-8c7f-f32e697dba4a:backup:91193c4f-45f0-949c-13e232f14c9e* backupCrn.

```
cdp.sh --no-verify-tls --endpoint-url https://cle-cdp.apps.drs131-1.vpc.
cloudera.com --no-verify-tls --form-factor private drscp describe-backup
   --backup-crn crn:altus:drs:us-west-1:18be-4c75-8c7f-f32e697dba4a:backup
:91193c4f-45f0-949c-13e232f14c9e
```

c) List all the backups using the list-backup CDP CLI option.

The following sample snippet output lists all the available backups.

```
cdp.sh --no-verify-tls --endpoint-url https://cle-cdp.apps.lhq-lp1-1.vpc
.cloudera.com --no-verify-tls --form-factor private drscp list-backup
```

d) Restore a specific backup, using its CRN, with the restore-backup CDP CLI option.

The following sample snippet restores the backup of *crn:altus:drs:us-west-1:88d84e3c-4c3e-9903-6c388a689690:backup:aebe-96d7-b79d10b64183* CRN.

```
cdp.sh --form-factor private --no-verify-tls --endpoint-url https://cle-ocp1.apps.srd-os-01.kcloud.cloudera.com drscp restore-backup --backup-cr n crn:altus:drs:us-west-1:88d84e3c-4c3e-9903-6c388a689690:backup:aebe-96 d7-b79d10b64183
```

e) Track the current status of the specified restoreCrn (restore event) using the describe-restore CDP CLI option. The following sample snippet output shows the current status of the *crn:altus:drs:us-west-1:a70c917a-4be8-927c-d36f3f7db2de:restore:c3b34532-4391-b62d-3f471fae5a40* restoreCrn:

```
cdp.sh --form-factor private --no-verify-tls --endpoint-url https://cle-
cp1.apps.srd-os-01.kcloud.cloudera.com drscp describe-restore --restore-
crn crn:altus:drs:us-west-1:a70c917a-4be8-927c-d36f3f7db2de:restore:c
3b34532-4391-b62d-3f471fae5a40
```

For information about all the available CDP CLI options to backup and restore Control Plane and CDW, see CDP CLI options for Control Plane namespaces and CDP CLI options for Cloudera Data Warehouse (CDW).

To set up a CDP client to run the CDP CLI commands, see CDP Private Cloud CLI.

DRS automatic backups (technical preview)

The DRS service can take automatic backups of the Control Plane and Cloudera Data Engineering (CDE) namespaces in the compute cluster of ECS. You can configure the schedule for periodic backup.



Important: The automatic backups of Control Plane and DEX namespaces is a technical preview feature and is not recommended for production deployments. Cloudera recommends that you try this feature in development or test environments. To enable this feature, contact your Cloudera account team.

You can choose one of the following storage options for DRS automatic backups in ECS:

External storage

ECS uses Longhorn as the underlying storage provisioner. In Longhorn, you can store snapshots externally using an S3 compatible storage such as Ozone or NFS v4. Cloudera recommends that you use external storage for automatic backups in ECS.

In-cluster storage

You can use Longhorn in-cluster storage. It is recommended that you use this option only if the external storage option is not available.

You can initiate the DRS automatic backups using the updateAutoBackupPolicy CDP CLI command. Alternatively, you can edit the automatic-backup (a Kubernetes cron job) to initiate the DRS automatic backups.

Configuring external storage in ECS for DRS automatic backups

Before you initiate DRS automatic backups, you must ensure that the prerequisite activities are completed along with the required additional configuration for external storage in Longhorn.

Before you begin



Important: Ensure that your deployment supports DRS automatic backups. The automatic backups of Control Plane and DEX namespaces is a technical preview feature and is not recommended for production deployments. Cloudera recommends that you try this feature in development or test environments. To enable this feature, contact your Cloudera account team.

Procedure

- 1. Complete the following prerequisites:
 - a) Ensure that the following requirements are met depending on the storage you choose for DRS automatic backups:
 - An S3 compatible storage, such as Ozone, must be available in the base cluster. You must have the
 required access key and secret to the storage, and the provisioned bucket must have a minimum of 5 TB
 storage space.
 - An NFS v4 storage must have a minimum of 5 TB of free space.
 - b) You must have SSH access to the base cluster node.



Note: If you are using Ozone storage, ensure that you have the SSH access to the base cluster running the S3 gateway service.

c) You must have SSH access to the ECS master node.



Tip: After you log into the terminal of the master node, run the following commands to access kubectl utility:

- 1. export PATH="/var/lib/rancher/rke2/bin:/opt/cloudera/parcels/ECS/
 docker/:\$PATH"
- 2. export KUBECONFIG=/etc/rancher/rke2/rke2.yaml
- **2.** Perform the following steps to change the default volume snapshot class value from *snap* (this value saves snapshots in the in-cluster storage in Longhorn) to *bak* (this value saves snapshots in the external storage in Longhorn):
 - a) Run the kubectl edit vsclass longhorn kubectl command.
 - b) Change the type parameter to bak as shown in the following sample snippet:

```
apiVersion: snapshot.storage.k8s.io/v1
deletionPolicy: Delete
driver: driver.longhorn.io
kind: VolumeSnapshotClass
metadata:
```

name: longhorn
parameters:
type: bak

- **3.** Complete the following steps if you are using Ozone S3 storage in Longhorn:
 - a) Run the scp root@[***BASE CLUSTER HOST***]:/var/lib/cloudera-scm-agent/agent-cert/cm-auto-global_cacerts.pem command to obtain the TLS certificate for Ozone.
 - DRS uses this certificate to communicate with the S3 gateway service using HTTPS.
 - b) Create a secret that Longhorn can use for S3 access. To accomplish this task, you must have the S3 access key, S3 secret, S3 endpoint, and S3 certificate for Ozone storage. You must also enable a virtual host to use the S3 compatible endpoint (Ozone).

The following sample snippet shows the kubectl command to create a secret:

```
kubectl create secret generic ozone-secret
--from-literal=AWS_ACCESS_KEY_ID=s3g/drs1-1.drs1.root.hwx.site@ROOT.HW
X.SITE
--from-literal=AWS_SECRET_ACCESS_KEY=9d9e46cc77bb510821f0dbc42c584a8b
7482b51dec9d3eb63c
--from-literal=AWS_ENDPOINTS=https://drs1.root.hwx.site:9879/longhorn
--from-literal=VIRTUAL_HOSTED_STYLE=true --from-file=AWS_CERT=cm-auto-g
lobal_cacerts.pem
-n longhorn-system
```

For more information, see Longhorn documentation.



Tip: Longhorn's URL is built using a combination of the AWS_ENDPOINTS value and the S3 virtual path as explained in Step 5.

4. Run the kubectl edit deploy cdp-release-thunderhead-drsprovider -n cdp-drs command, and set the TAKE_PVC_CLONE environment value to false.

This step ensures that the backups do not create a PVC clone for external snapshot.

By default, Longhorn configuration is set to in-cluster storage and this storage requires a PVC copy to perform the DRS restore operation (DRS uses CSI snapshot technology). Therefore, to use the external storage, you must configure the volume snapshot class to *bak* and then configure the TAKE_PVC_CLONE environment value to false.

- **5.** Configure the volume for NFS storage or bucket for Ozone S3 (on the Setting General page) in the Longhorn UI to save the backups.
 - a) Enter the nfs://... URL in the **Backup Target** field if you are using NFS storage.
 - b) Enter the required values in the following fields if you are using Ozone S3 storage:
 - s3://[***BUCKET***]@[***DUMMYREGION***]/ URL in the **Backup Target** field. For example, s3:// drs1-1@cdp/.
 - [***SECRET THAT YOU GENERATED IN STEP 3B***] in the **Backup Target Credential Secret**. For example, ozone-secret.

The s3://[***BUCKET***]@[***DUMMYREGION***]/ URL is a virtual S3 URL that you can create using the original Ozone S3 URL, where,

- bucket is the hostname. Longhorn prefixes the AWS_ENDPOINTS to the bucket value. For example, the sample snippet in Step 3 shows the hostname value as drs1-1.drs1.root.hwx.site. In this instance, drs1-1 is the bucket name and the rest of the hostname drs1.root.hwx.site is the AWS_ENDPOINTS hostname.
- dummyregion can be any value and is not used.

Troubleshooting: To verify whether Longhorn successfully registered the Ozone S3 credential secret, click the Backup page. No errors must appear on the page.

If any error or message appears about the secret and the certificate having newlines or space, run the kubectl edit lhs backup-target-credential-secret -n longhorn-system command and set the value to the secret you created in Step 3b.

What to do next

Initiate the DRS automatic backups using the updateAutoBackupPolicy CDP CLI command. Alternatively, you can edit the "automatic-backup" (a Kubernetes cron job) to initiate the DRS automatic backups.

Initiating DRS automatic backups

After you configure the external storage in ECS, you can initiate the DRS automatic backups using the "updateAutoBackupPolicy" CDP CLI command. Alternatively, you can edit the "automatic-backup" (a Kubernetes cron job) to initiate the DRS automatic backups.

Before you begin

The preferred method to initiate the DRS automatic backups is to use the updateAutoBackupPolicy CDP CLI command in the CDP client. For more information about DRS CDP CLI commands, see CLI reference for using DRS on Control Plane.

About this task

The following steps show an alternate method to initiate DRS automatic backups using kubectl commands.

Procedure

- 1. Run the kubectl edit cj automatic-backup -n cdp-drs command.
- 2. Configure the ENABLED environment variable to true to enable automatic backups, configure the namespaces (if they are not configured), and then configure the backup retain count to take backups on an hourly, daily, or weekly basis. You can also choose a combination of two or more periods to take backups. Save the cron job.

The backup retain count determines the number of backup instances to generate.



Important: If you do not want to set a backup for a particular period, ensure that the count is set to 0. This is because the retain count is set to 1 (minimum backup retain count) by default.

DRS generates n+1 backups by default where n is the backup retain count. Therefore, the minimum number of backups at any point in time is 2 by default. For example, if you set the HOURLY_COUNT parameter to 2, three instances are generated; therefore, two backups are taken every hour. If you set the WEEKLY_COUNT parameter to 0, no instances are created and no backups are generated.

The following sample snippet shows the environment variables required for DRS automatic backups:

```
env:
    name: ENABLED
value: "true"
    name: HOURLY_COUNT
value: "1"
    name: DAILY_COUNT
value: "1"
    name: WEEKLY_COUNT
value: "1"
```

Results

By default, Kubernetes initiates the first automatic backup within 30 minutes after the backup policy creation is complete.



Tip: You can configure the cron schedule environment variable (using kubectl commands or updateAutoBackupPolicy CDP CLI command) to control the next job run.

The cron schedule uses the [***MINUTE FROM 0-59***] [***HOUR FROM 0-23***] [***DAY OF THE MONTH FROM 1-31***] [***MONTH FROM 1-12***] [***DAY OF THE WEEK FROM 0-7 WHERE 7 IS FOR SUNDAY***] cron syntax.

For example, if you configure schedule: '*/40 * * * *', the backup runs after 40 minutes.

What to do next

Backup instances, depending on the chosen schedules, are generated and appear on the CDP Private Cloud Data Services Management Console Dashboard Backup Overview View Details Backup and Restore Manager Backups tab. The instance name is auto-generated. Click the backup instance to view more details.

CLI reference for using DRS on Control Plane

You can use the Data Recovery Service (DRS) CDP CLI commands to backup and restore resources and data in the Control Plane of CDP Private Cloud Data Services.

The following table provides the CDP CLI options to backup and restore the Control Plane:

CDP CLI options	Description
create-backup	Creates a backup and archives it as a ZIP file on the same cluster.
	The item-name is optional for Control Plane data recovery service.
	You can provide a unique backup name, so that you can identify the backup easily during restore.
	When you run this command, the service initiates the backup process and returns the assigned ID or backupCrn for the backup.
delete-backup	Deletes the specified backup (backupCrn) permanently.
	Important: Alerts are not generated when you run this command.
describeAutoBackupPolicy	Describes the automatic backup policy. You can view the DRS automatic backup and the selected periods (hourly, daily, weekly) and its backup retain count. For each control plane, only one automatic-backup cron job can be scheduled in the <i>drs</i> namespace.
describe-backup	Shows the progress of the current status of the specified backupCrn (backup event).
describe-restore	Shows the progress of the current status of the specified restoreCrn (restore event).
get-logs	Returns logs about the specified backup, restore, or delete job and automatically creates a diagnostic bundle.
	You can download the bundle to your machine to analyze an issue or share it with Cloudera Support for further troubleshooting.
list-backup-entities	Lists the entities that you can backup, which includes the Control Plane namespace and its corresponding vault namespace (if embedded).
list-backups	Lists the successful backup jobs of backupCrn.
	You can filter the backup jobs using the NOT_STARTED; IN_PROGRESS; COMPLETED; PARTIALLY_FAILED; and FAILED job states.
list-restores	Lists the past restore events.

CDP CLI options	Description
restore-backup	Restores the backup of the specified backupCrn. During the restore event, the existing Kubernetes resources and data are deleted and then recreated using the information in the backup. When you run the command, the service initiates the restore event and returns a restoreCrn value.
updateAutoBackupPolicy	Creates or updates the current existing automatic backup policy. The
apane. Into Buckupi Oney	automatic backup cron job starts with the default backup policy first.
	You can enable or disable DRS automatic backups. You can also add, delete, and update the period (hourly, daily, weekly), and its backup retain count.

For more information about the DRS CDP CLI options for Control Plane, see CDP CLI drscp.

Related Information
CDP CLI drscp

Troubleshooting Backup and Restore Manager

The troubleshooting scenarios in this topic help you to troubleshoot issues that might appear for DRS in the Control Plane. The "Backup and Restore Manager" in CDP Private Cloud Data Services Management Console leverages the Data Recovery Service (DRS) capabilities to backup and restore Kubernetes namespaces and resources.

CDP Control Plane UI or the Backup and Restore Manager becomes inaccessible after a failed restore event?

Problem

CDP Control Plane UI does not come up or the Backup and Restore Manager (or drscp options) becomes inaccessible after a failed restore event.

Cause

Sometimes, some configurations take more time to restore. For example, in a shared cluster (OCP) that is heavily loaded, the restore event might surpass the set timeout limit. In this scenario, you can either wait or rerun the restore event again.



Tip: Run the restore event for such scenarios during non-peak hours.

Solution

You can perform one of the following steps after a failed restore event:

- Wait for a minimum of 15 minutes. This might resolve the issue automatically if the issue was caused due to timeout. You can verify this in the logs.
- Run restore again. This might resolve the issue if it was temporary such as, restore event during cluster maintenance.

If the Control Plane is not restored successfully even after you follow the steps, contact Cloudera Support for further assistance.

Timeout error appears in Backup and Restore Manager

Problem

A timeout error appears in the Backup and Restore Manager (or drscp options) during a restore event.

Solution

When the restore event crosses the time set in the POD_CREATION_TIMEOUT environment property of the cdp-release-thunderhead-drsprovider deployment in the drs namespace, a timeout error appears. By default, the property is set to 900 seconds. In this scenario, you must manually verify whether the pods are up or not.

Timeout error during backup of OCP clusters

Problem

"The execution of the sync command has timed out" error appears during a backup event for OCP clusters.

Cause

This scenario is observed when the cluster is heavily used and the backup event is initiated during peak hours.

Solution

You can restart the nodes, this causes the disk to unmount and forces the operating system to write any data in its cache to the disk. After the restart is complete, initiate another backup. If any warnings appear, scrutinize to verify whether there are any dire warnings, otherwise the generated backup is safe to use. The only drawback in this scenario is the downtime impact, that is the time taken to back up the OCP clusters is longer than usual. Therefore, it is recommended that you back up the clusters during non-peak hours.

If the sync errors continue to appear, contact your IT department to check whether there is an issue with the storage infrastructure which might be preventing the sync command from completing on time.

Stale configurations in Cloudera Manager after a restore event

Cause

This scenario appears when you take a backup of the CDP Private Cloud Data Services Control Plane, upgrade Data Services, and then perform a restore. During the upgrade process, new parcels are activated and configurations in Cloudera Manager might have changed.

Solution

It is recommended that you restart Cloudera Manager after the upgrade process is complete and then initiate the restore event.

Restore event for an environment backup fails with an exception

Problem

When you delete an environment after the backup event, the restore operation for the environment fails and the Not able to fetch details from Cluster:... exception appears.

Cause

During the environment creation process, the environment service creates an internal Cloudera Manager user with Full Administrator role. The username is stored in the Control Plane database, and the password is stored in the vault. When you delete an environment, the internal Cloudera Manager user gets deleted. The exception appears only if the password is no longer valid or might be missing. One of the reasons why the password might go missing is that while fixing a vault corruption, the vault might have been rebuilt without fixing the Cloudera Manager credentials.

Solution

Procedure



Note: The following commands use the cdp-embedded-db-0 environment:

Get the internal Cloudera Manager username using the following commands to determine whether the credential is valid.

- **a.** Login into the environment using the kubectl exec -it cdp-embedded-db-0 -n [***CONTROL PLANE NAMESPACE***] psql command.
- **b.** Connect to the environment database using the \c db-env; command.
- **c.** Run the following SQL query in the cdp-embedded-db-0 pod:

```
SELECT e.environment_crn, c.value FROM environments e JOIN configs c ON e.environment_crn = c.environment_crn WHERE e.environment_name = '[***YOUR ENV NAME***]' AND c attr = 'cmUser';
```

Sample output:

The __cloudera_internal_user__ test-saml2-env-1-798414fe-faa6-43e1-ac9c-75c4d33ec294 is the internal Cloudera Manager username in the sample output.

- **2.** Get the internal Cloudera Manager password using the following commands:
 - a) Run the following commands to get the root token for the embedded vault:
 - 1. If you are using OCP:

```
$ kubectl get secret vault-unseal-key -n [***VAULT-NAMESPACE***] -o jsonpath="{. data.init\.json}" | base64 -d {"keys":["[***VALUE***]"],"keys_base64":["[***value***]="],"recovery_keys":null,"recovery_keys_base64":null,"root_token":"[***VALUE***]"} command returns the vault root token.
```

- 2. If you are using ECS:
 - a. [root@cm_server_db_host ~]# psql -U cm cm
 - select * from CONFIGS where attr like '%vault_root%';

Sample output:

1546337327	vault_root	hvs.SvIrIhhffYEmVPEWN3TSEzks	15463371
54 0	NONE		

The hvs.SvIrIhhffYEmVPEWN3TSEzks value in the above sample output is the vault token.

- b) kubectl exec -it vault-0 -n [***VAULT_NAMESPACE***] /bin/sh
- c) export VAULT TOKEN=[***VAULT ROOT TOKEN***]
- d) ~ \$ vault secrets list -detailed -tls-skip-verify

Sample output:

```
Plugin
Path
                            Accessor
                                                    Default TTL
   Force No Cache Replication Seal Wrap External Entropy Access
    Options
                      Description
         UUID
                                                   Version
                                                             Running Vers
ion
             Running SHA256 Deprecation Status
cubbyhole/ cubbyhole cubbyhole_35ff7854 n/a false local false false map[] per-token private secret storage
                                                                    n/a
                      per-token private secret storage
    map[]
       f2fa15ec-49-cea2-88f6-e6807c30fba3 n/a v1.13.1+builti
n.vault n/a n/a identity/ identity identity_b7aa2294 system false replicated false false
                                                                    system
                       identity store
     map[]
          17990faa-e0-727a-92a3-aaaa1ff43393 n/a v1.13.1+built
          n/a n/a
kv kv_2ba3b77c system
e replicated false false
in.vault
kv/
                                                                    system
     false
           map[version:2] key/value secret storage
                98b14495-b6-6958-04bc-1ca7c55d4590 n/a
     14.2+builtin n/a supported et/ kv kv_218f4379 system false replicated false false map[version:2] key/value secret storage
                    n/a
v0.14.2+builtin
secret/ kv
                                                                    system
    false
          06371963-e6-56c1-7ab3-d6c438720dbf n/a v0.14.2+built
           n/a supported
system system_46e657a4 n/a
replicated true false
in
sys/
                                                                    n/a
   false
                     system endpoints used for control, policy and debu
    map[]
gging 8ca5d96f-a45e-155a-cfc1-25a56b6a0de5 n/a v1.13.1+buil
tin.vault
           n/a
```

In this command output, kv/ is the secret path.

e) ~ \$ vault kv list -tls-skip-verify kv

Sample output:

```
Keys
----
[***CONTROL PLANE NAMESPACE***]
```

f) ~ \$ vault kv list -tls-skip-verify kv/[***CONTROL PLANE NAMESPACE***]

Sample output:

```
Keys
----
data/
liftie/
```

test

g) ~ \$ vault kv list -tls-skip-verify kv/<[***CONTROL PLANE NAMESPACE***]/data Sample output:

```
Keys
----
[***ENV NAME1***]
[***ENV NAME2***]
```

Identify the environment for which the exception appeared.

h) ~ \$ vault kv list -tls-skip-verify NAME WITH THE EXCEPTION***] kv/[***CONTROL PLANE NAMESPACE***]/[***ENTER THE ENV

Sample output:

```
Keys
----
[***RANDOM UUID***]
```

i) ~ \$ vault kv list -tls-skip-verify kv/[***CONTROL PLANE NAMESPACE***]/[***ENTER THE ENV NAME WITH THE EXCEPTION***]/[***RANDOM UUID***]

Sample output:

```
Keys
----
cmPassword
dockerConfigJson
kubeconfig
```

j) ~ \$ vault kv get -tls-skip-verify kv/[***CONTROL PLANE NAMESPACE***]/[***ENTER THE ENV NAME WITH THE EXCEPTION***]/[***RANDOM UUID***]/cmPassword

Sample output:

```
kv/[***CONTROL PLANE NAMESPACE***]/[***ENV NAME***]/[***RANDOM UUID***]/
cmPassword
===== Metadata ======
               Value
Key
created_time 2023-11-15T04:32:36.477837897Z
custom_metadata <nil>
deletion_time n/a
              false
destroyed
version
==== Data ====
Key
   Value
value
       ae4cff8a-fcee-48e9-b381-4a16e883694a88c8d2
```

The value is the cmPassword (Cloudera Manager password).

3. Log into Cloudera Manager using the username (cloudera_internal_user) and password (cmPassword) that you obtained in the previous steps.

- 4. Run the following commands as shown to regenerate the internal Cloudera Manager credentials in bash:
 - a) [root@user ~]# uuidgen command creates the first universally unique identifier (UUID) which you use in the Cloudera Manager username.

Sample output:

```
dc7c7dd7-5a58-497a-a1d1-46cd
```

b) [root@user ~]# uuidgen command creates another universally unique identifier (UUID) which is the Cloudera Manager password.

Sample output:

```
9a863dc4-be61-430f-ac87-a4eba0
```

5. Assemble the new Cloudera Manager username using the information from the previous commands in the "__c loudera_internal_user__" + [***ENTER THE ENV NAME WITH THE EXCEPTION***] + "-" + [***FIRST_UUID***] format.

For example, __cloudera_internal_user__cldrienv1-dc7c7dd7-5a58-497a-a1d1-46cd. In this assembled Cloudera Manager username, the prefix __cloudera_internal_user__ is followed by a string that contains the name of the environment with the exception cldrienv1 and the generated UUID dc7c7dd7-5a58-497a-a1d1-46cd separated by "_"

The new Cloudera Manager password is the second UUID. For example, 9a863dc4-be61-430f-ac87-a4eba0

- **6.** Go to the Cloudera Manager Support API Explorer UsersResource POST /users REST API, and perform the following steps:
 - a) Click Try it out, and substitute the Cloudera Manager username and password in the following JSON string:

b) Copy the JSON string into the REQUEST BODY, and click Execute.

You get 200 response code.

- 7. Verify whether you can use the username and password to log into Cloudera Manager.
- 8. Replace the stale Cloudera Manager user with the new username with the following commands:
 - a) kubectl exec -it cdp-embedded-db-0 -n [***CONTROL PLANE NAMESPACE***] psql
 - b) \c db-env;
 - c) Run the following SQL queries in the cdp-embedded-db-0 pod:
 - 1. SELECT e.environment_crn, c.value FROM environments e JOIN configs c ON e.environment_crn c.environment_crn WHERE e.environment_name = '[***YOUR ENV NAME***]' AND c.attr = 'cmUser';
 - **2.** UPDATE configs SET value='[***NEW CLOUDERA MANAGER INTERNAL USER***]' WH ERE environment_crn='[***ENVIRONMENT CRN OF ENV WITH THE EXCEPTION***]' AND attr='cmUser'; command replaces the old Cloudera Manager username.

- **9.** Replace the stale Cloudera Manager password with the new password:
 - a) Run the steps in Step 2 to find the Cloudera Manager user password credential path in the vault which should be in kv/[***CONTROL PLANE NAMESPACE***]/[***ENV-NAME***]/[***RANDOM UUID***]/cmP assword format.

 - c) Verify whether the cmPassword is changed using the \$\ vault kv get -tls-skip-verify kv/[***CONTROL PLANE NAMESPACE***]/[***ENV NAME WITH THE EXCEPTION***]/[***RANDOM UUID***]/cmP assword command.

Using DRS with CDW

You can back up and restore Kubernetes namespaces behind Cloudera Data Warehouse (CDW) entities (for example, Database Catalogs, Virtual Warehouses) on demand using the Data Recovery Service (DRS). CDW leverages DRS and provides CDP CLI endpoints which you can use to create and restore backups for CDW namespaces to back up CDW metadata and configurations such as Kubernetes objects, persistent volumes, autoscaling configuration, and so on.

The following limitations apply for CDW data service if you are on Embedded Container Service (ECS) or using an embedded database on Red Hat OpenShift Container Platform:

- The embedded database that CDW uses is part of the Control Plane. You cannot back up only CDW-related
 entities from the embedded database using the dw create-backup command. You must take a backup of the
 Control Plane service.
- You must restore the entire Control Plane configurations to restore configurations stored in the CDW database.
 This recreates the Control Plane namespace.

List of data recovery sub-commands for CDW

The following table lists the commands and CLI endpoints for backing up and restoring Kubernetes namespaces behind CDW entities:

DRS sub-commands for CDW	Description
create-backup	Creates an on-demand backup for the Data Warehouse including Kubernetes objects, persistent volumes, and so on. Backup requests are processed asynchronously and instantaneously.
delete-backup	Deletes an existing Data Warehouse backup. The call returns immediately. It returns a delete CRN, which is the deletion process identifier.
describe-backup	Returns the description of an existing Data Warehouse backup.
restore-backup	Restores the state of the Data Warehouse from an existing backup. It returns a restore CRN, which is the identifier of the restoration process.
describe-restore	Returns the description of the Data Warehouse restore operation.
list-backup-entities	Lists potential backup entities associated with the Data Warehouse.
list-backups	Lists backups associated with the Data Warehouse.
list-restores	Lists restores associated with the Data Warehouse.
get-logs	Returns the job logs corresponding to the specified CRN.

Related Information

CDP CLI reference of DRS commands for CDW