

Migrating a Deployment to a New Set of Hosts

Date published: 2020-02-28

Date modified:

The Cloudera logo is displayed in a bold, orange, sans-serif font. The word "CLOUDERA" is written in all caps, with a stylized 'E' that has a horizontal bar extending to the right.

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

Adding and Removing Worker Hosts.....	4
Migrating a CSD Deployment.....	4
Add and Set Up the New Hosts.....	4
Copy the JDK to the new host.....	4
Copy the DNS Nameserver to the new host.....	4
Copy the Kerberos Configurations.....	4
Stop the CDSW Service.....	4
Backup Application Data.....	5
Delete CDSW Roles from Existing Hosts.....	5
Move Backup to the New Master.....	5
Update DNS Records for the New Master.....	5
Add Role Instances for the New Hosts.....	5
Run the Prepare Node command on the New Hosts.....	6
Start the CDSW Service.....	6
Migrating an RPM Deployment.....	7
Add and Set Up the New Hosts.....	7
Copy the JDK to the new host.....	7
Copy the DNS Nameserver to the New Host.....	7
Copy the Kerberos Configurations.....	7
Stop Cloudera Data Science Workbench.....	7
Backup Application Data.....	8
Remove Cloudera Data Science Workbench from Existing Hosts.....	8
Move Backup to New Master.....	8
Update DNS Records for the New Master.....	8
Install Cloudera Data Science Workbench on New Master Host.....	8

Adding and Removing Worker Hosts

This section describes how add and remove worker hosts from Cloudera Data Science Workbench.

Migrating a CSD Deployment

This section describes how to migrate a CSD-based Cloudera Data Science Workbench service to a new set of gateway hosts.

Add and Set Up the New Hosts

Add and set up the new hosts for migrating a CSD deployment.

Procedure

1. [Add new hosts to your cluster as needed.](#)

Make sure they are gateway hosts that have been assigned gateway roles for HDFS, YARN, and Spark 2. Do not run any other services on these hosts.

2. Set up the new hosts as per the Cloudera Data Science Workbench hardware requirements listed [here](#).
 - [Disable Untrusted SSH Access](#) on the new hosts.
 - [Configure Block Devices](#) on the new hosts.

Copy the JDK to the new host

Provides information on how to copy the JDK to the new host.

Procedure

Copy the /usr/java directory to the new host.

Copy the DNS Nameserver to the new host

Provides information on how to copy the DNS nameserver to the new host.

Procedure

Copy the /etc/resolv.conf file to the new host.

Copy the Kerberos Configurations

Provides information on how to copy the Kerberos configurations.

Procedure

Copy the /etc/jkr5.conf file to the new host.


Stop the CDSW Service

Use Cloudera Manager to stop all roles of the CDSW service.

About this task

On Cloudera Data Science Workbench 1.4.0 (and lower), do not stop or restart Cloudera Data Science Workbench without using the [cdsw_protect_stop_restart.sh](#) script. This is to help avoid the data loss issue detailed in TSB-346.

Procedure

1. Log into the Cloudera Manager Admin Console.
2. On the Home Status tab, click  to the right of the CDSW service and select Stop from the dropdown.
3. Confirm your choice on the next screen. When you see a Finished status, the action is complete.

Backup Application Data

In Cloudera Data Science Workbench all stateful data is stored on the master host at `/var/lib/cdsw`. Backup the contents of this directory before you begin the migration process.

Procedure

1. After stopping CDSW, and before running the following tar command, wait 2-5 minutes (depending on your disk speed) to ensure that all data from CDSW is successfully written to the disks. Otherwise the tar command may not capture all recent changes.
2. To create the backup, run the following command on the master host:

```
tar cvzf cdsw.tar.gz /var/lib/cdsw/*
```

Delete CDSW Roles from Existing Hosts

This provides the steps for deleting CDSW roles from an existing host.

Procedure

1. Log into the Cloudera Manager Admin Console.
2. Go to the CDSW service.
3. Click the Instances tab.
4. Select all the role instances.
5. Select Actions for Selected > Delete. Click Delete to confirm the deletion.

Move Backup to the New Master

This provides steps to move the backup to the new master.

Procedure

1. Copy the backup taken previously to the host that will be the new Cloudera Data Science Workbench master.
2. Unpack the contents of the backup into `/var/lib/cdsw`.

```
tar xvzf cdsw.tar.gz -C /var/lib/cdsw
```

Update DNS Records for the New Master

Update your DNS records with the IP address for the new master host.

Add Role Instances for the New Hosts

This provides steps to add role instances for the new host.

Procedure

1. Log into the Cloudera Manager Admin Console.
2. Go to the CDSW service.
3. Click the Instances tab.

4. Click Add Role Instances. Assign the Cloudera Data Science Workbench Master, Application, and Docker Daemon roles to the new master host. If you want to configure worker hosts, assign the Cloudera Data Science Workbench Worker and Docker Daemon roles to the new workers.



Important: Do not assign the Master and Worker roles to the same host. Even if you only have one host for Cloudera Data Science Workbench, the Master can automatically perform the functions of a Worker host as needed.

5. Click Continue. On the Review Changes page, review the configuration changes to be applied. The wizard finishes by performing any actions necessary to add the new role instances.

Do not start the new roles at this point. You must run the Prepare Node command as described in the next step before the roles are started.

Run the Prepare Node command on the New Hosts

This provides information on how to run the prepare command on the new hosts.

Procedure

1. The new hosts must have the following packages installed on it.

```
nfs-utils
libseccomp
lvm2
bridge-utils
libtool-ltdl
iptables
rsync
policycoreutils-python
selinux-policy-base
selinux-policy-targeted
ntp
ebtables
bind-utils
openssl
e2fsprogs
redhat-lsb-core
conntrack-tool
bash
curl
```

You can either manually install these packages now, or, allow Cloudera Manager to install them as part of the Prepare Node command later in this step.

If you choose the latter, make sure that Cloudera Manager has the permissions needed to install the required packages. To do so, go to the CDSW service and click Configuration. Search for the Install Required Packages property and make sure it is enabled.


2. Run the Prepare Node command on the new hosts.
 - a) Go to the CDSW service.
 - b) Click Instances.
 - c) Select all the role instances.
 - d) Select Actions for Selected > Prepare Node. This will install the required set of packages on all the new hosts.

Start the CDSW Service

This provides the steps to start the CDSW service.

Procedure

1. Log into the Cloudera Manager Admin Console.

2. On the Home Status tab, click  to the right of the CDSW service and select Start from the dropdown.
3. Confirm your choice on the next screen. When you see a Finished status, the action is complete.

Migrating an RPM Deployment

This section describes how to migrate an RPM-based Cloudera Data Science Workbench service to a new set of gateway hosts.

Add and Set Up the New Hosts

Provides the steps to add and set up the new hosts.

Procedure

1. [Add new hosts to your cluster as needed](#). Make sure they are gateway hosts that have been assigned gateway roles for HDFS, YARN, and Spark 2. Do not run any other services on these hosts.
2. Set up the new hosts as per the Cloudera Data Science Workbench hardware requirements listed [here](#).
 - [Disable Untrusted SSH Access](#) on the new hosts.
 - [Configure Block Devices](#) on the new hosts.

Copy the JDK to the new host

Provides the steps to copy the JDK to the new host.

Procedure

Copy the /usr/java directory to the new host.

Copy the DNS Nameserver to the New Host

Provides the steps to copy the DNS nameserver to the new host.

Procedure

Copy the /etc/resolv.conf file to the new host.

Copy the Kerberos Configurations

Provides the steps to copy the Kerberos configurations.

Procedure

Copy the /etc/jkr5.conf file to the new host.

Stop Cloudera Data Science Workbench

Provides the steps to stop the Cloudera Data Science Workbench service.

About this task

On Cloudera Data Science Workbench 1.4.0 (and lower), do not stop or restart Cloudera Data Science Workbench without using the [cdsw_protect_stop_restart.sh](#) script. This is to help avoid the data loss issue detailed in TSB-346.

Procedure

Run the following command on the master host to stop Cloudera Data Science Workbench.

```
cdsw stop
```

Backup Application Data

Provides the steps to back up the application data.

About this task

In Cloudera Data Science Workbench all stateful data is stored on the master host at `/var/lib/cdsw`. Backup the contents of this directory before you begin the migration process.

Procedure

1. After stopping CDSW, and before running the following tar command, wait 2-5 minutes (depending on your disk speed) to ensure that all data from CDSW is successfully written to the disks. Otherwise the tar command may not capture all recent changes.
2. To create the backup, run the following command on the master host:

```
tar cvzf cdsw.tar.gz /var/lib/cdsw/*
```

Remove Cloudera Data Science Workbench from Existing Hosts

Provides the steps to remove Cloudera Data Science Workbench from the existing hosts.

Procedure

Run the following commands on the existing master and any worker hosts you want to migrate:

```
cdsw stop
yum remove cloudera-data-science-workbench
```

Move Backup to New Master

Provides the steps to move the backup to a new master.

Procedure

Copy the backup taken in the previous step to the host that will be the new Cloudera Data Science Workbench master. Unpack the contents of the backup into `/var/lib/cdsw`.

```
tar xvzf cdsw.tar.gz -C /var/lib/cdsw
```

Update DNS Records for the New Master

Update your DNS records with the IP address for the new master host.

Install Cloudera Data Science Workbench on New Master Host

Provides information on how to install Cloudera Data Science Workbench on the new master host.

Procedure

For instructions, see [Installing Cloudera Data Science Workbench 1.8.1 using Packages](#).