

CDP Private Cloud Base 7

CDP Private Cloud Base Edition Release Guide

Date published: 2019-11-22

Date modified: 2024-07-19

CLOUDEXERA

<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

CDP Private Cloud Base Release Guide.....	5
CDP Release Notes.....	5
Version and Download Information.....	6
Cloudera Manager Version Information.....	6
Cloudera Manager Download Information.....	7
Cloudera Runtime Version Information.....	24
Cloudera Runtime Download Information.....	25
CDP Private Cloud Base Trial Download Information.....	27
Product Compatibility Matrices.....	28
Replication Manager on CDP Private Cloud Base.....	28
Cloudera Manager with Runtime, CDH, and Data Services.....	31
Product Compatibility Matrices for KMS and Encryption Products.....	36
Product Compatibility Matrix for Ranger KMS.....	36
Product Compatibility Matrix for Navigator Encrypt.....	37
Product Compatibility Matrix for KTS and Key HSM.....	38
Product Compatibility Matrix for HSM Support.....	40
Changes to CDH and HDP Components in CDP Private Cloud Base.....	41
Updated CDH Components.....	41
Updated HDP Components.....	44
HDP Core component version changes.....	45
Changes to Ambari and HDP services.....	45
Assessing the Impact of Apache Hive.....	46
Apache Hive key features.....	47
Apache Hive 3 architectural overview.....	48
Key semantic changes and workarounds.....	50
Casting timestamps.....	50
Casting invalid dates.....	50
Changing incompatible column types.....	51
Understanding CREATE TABLE behavior.....	51
Configuring legacy CREATE TABLE behavior.....	53
Handling table reference syntax.....	53
Add Backticks to Table References.....	54
Handling the Keyword APPLICATION.....	54
Dropping partitions.....	54
Handling output of greatest and least functions.....	55
Renaming tables.....	55
TRUNCATE TABLE on an external table.....	55
Hive unsupported interfaces and features.....	56

Supplemental Upgrade Topics.....	57
Configuring a Local Package Repository.....	57
Creating a Permanent Internal Repository.....	58
Creating a Temporary Internal Repository.....	59
Configuring Hosts to Use the Internal Repository.....	59
Configuring a Local Parcel Repository.....	60
Using an Internally Hosted Remote Parcel Repository.....	60
Using a Local Parcel Repository.....	62
Changes to CDH Hive Tables.....	63
Changes to HDP Hive tables.....	64
Transitioning from the Cloudera Manager Embedded PostgreSQL Database Server to an External PostgreSQL Database.....	65
Prerequisites.....	66
Identify Roles that Use the Embedded Database Server.....	66
Migrate Databases from the Embedded Database Server to the External PostgreSQL Database Server.....	68

CDP Private Cloud Base Release Guide

This guide contains release and download information for installers and administrators. It includes links to release notes as well as information about installation requirements, supported platforms, and version and download information.

CDP Release Notes

Links to the Cloudera Manager and Cloudera Runtime release notes for CDP Private Cloud Base and CDP Private Cloud Data Services.

Table 1: Release Note Links

CDP Version	Release Note Links
7.1.9 SP1 CDP Private Cloud Base	Cloudera Manager 7.11.3 Cumulative hotfix 7 Release Notes Cloudera Runtime 7.1.9 SP1 Release Notes
7.1.7 SP3 CDP Private Cloud Base	Cloudera Manager 7.11.3 Release Notes Cloudera Runtime 7.1.7 SP3 Release Notes
7.1.9 CDP Private Cloud Base	Cloudera Manager 7.11.3 Release Notes Cloudera Runtime 7.1.9 Release Notes
7.1.7 SP2 CDP Private Cloud Base	Cloudera Manager 7.6.7 Release Notes Cloudera Runtime 7.1.7 SP2 Release Notes
CDP Private Cloud Data Services 1.4.1	Cloudera Manager 7.8.1 Release Notes Cloudera Runtime 7.1.8 Release Notes
7.1.8 CDP Private Cloud Base	Cloudera Manager 7.7.1 Release Notes Cloudera Runtime 7.1.8 Release Notes
CDP Private Cloud Data Services 1.4.0	Cloudera Manager 7.6.5 Release Notes Cloudera Runtime 7.1.7 Release Notes
CDP Private Cloud Data Services 1.3.4	Cloudera Manager 7.5.5 Release Notes Cloudera Runtime 7.1.7 Release Notes
CDP Private Cloud Data Services 1.3.3	Cloudera Manager 7.5.4 Release Notes Cloudera Runtime 7.1.7 Release Notes
CDP Private Cloud Data Services 1.3.2	Cloudera Manager 7.5.4 Release Notes Cloudera Runtime 7.1.7 Release Notes
CDP Private Cloud Data Services 1.3.1	Cloudera Manager 7.5.1 Release Notes Cloudera Runtime 7.1.7 Release Notes
7.1.7 CDP Private Cloud Base	Cloudera Manager 7.6.1 (Service Pack 1 for Cloudera Manager 7.4.4) Release Notes Cloudera Manager 7.4.4 Release Notes Cloudera Runtime Release Notes

CDP Version	Release Note Links
7.1.6 CDP Private Cloud Base	Cloudera Manager 7.3.1 Release Notes Cloudera Runtime Release Notes
7.1.5 CDP Private Cloud Base	Cloudera Manager 7.2.4 Release Notes Cloudera Runtime 7.1.5 Release Notes
7.1.4 CDP Private Cloud Base	Cloudera Manager 7.1.4 Release Notes Cloudera Runtime Release Notes
7.1.3 CDP Private Cloud Base	Cloudera Manager 7.1.3 Release Notes Cloudera Runtime Release Notes
7.1.2 CDP Private Cloud Base	Cloudera Manager 7.1.2 Release Notes Cloudera Runtime Release Notes
7.1.1 CDP Private Cloud Base	Cloudera Manager 7.1.1 Release Notes Cloudera Runtime Release Notes
7.0.3 CDP Private Cloud Base	Cloudera Manager 7.0.3 Release Notes Cloudera Runtime Release Notes

Version and Download Information

The following topics describe the available versions and download locations for Cloudera Manager and Cloudera Runtime.

Cloudera Manager Version Information

You must choose the correct Cloudera Manager for your deployment. This page provides a reference of Cloudera Manager versions, their release dates, and important compatibility information.

Cloudera Manager 7.11.3 CHF7 is the current release of Cloudera Manager required for CDP Private Cloud Base version 7.1.9 SP1.

Cloudera Manager 7.11.3 CHF4 is the current release of Cloudera Manager required for CDP Private Cloud Base version 7.1.7 SP3.



Important:

You must install Python 3 on all hosts before installing or upgrading to Cloudera Manager 7.11.3. For more information, see the [Installing Python 3](#).

Release date: July 19, 2024

Previous releases:

- Cloudera Manager 7.11.3 CHF4 Release Date: April 05, 2024
- Cloudera Manager 7.11.3 Release Date: August 18, 2023
- Cloudera Manager 7.10.1 Release Date: June 13, 2023
- Cloudera Manager 7.9.5 Release Date: January 25, 2023
- Cloudera Manager 7.8.1 Release Date: November 18, 2022
- Cloudera Manager 7.7.3 Release Date: October 28, 2022
- Cloudera Manager 7.7.1 Release Date: August 30, 2022
- Cloudera Manager 7.6.5 Release Date: May 25, 2022

- Cloudera Manager 7.6.1 (SP1) Release Date: March 30, 2022
- Cloudera Manager 7.5.5 Release Date: April 13, 2022
- Cloudera Manager 7.5.4-20668437 Release Date: January 13, 2022
- Cloudera Manager 7.5.4 Release Date: November 8, 2021
- Cloudera Manager 7.5.1 Release Date: October 4, 2021
- Cloudera Manager 7.4.4 Release Date: August 5, 2021
- Cloudera Manager 7.3.1 Release Date: March 3, 2021
- Cloudera Manager 7.2.4 Release Date: November 30 2020
- Cloudera Manager 7.1.4 Release Date: October 13, 2020
- Cloudera Manager 7.1.3 Release Date: August 10, 2020
- Cloudera Manager 7.1.2 Release Date: July 13, 2020
- Cloudera Manager 7.1.1 Release Date: May 22, 2020
- Cloudera Manager 7.0.3 Release Date: November 22, 2019

Cloudera Manager Download Information

Important: Access to Cloudera Manager binaries for production purposes requires authentication. To access the binaries at the locations below, you must first have an active subscription agreement and obtain a license key file along with the required authentication credentials (username and password).

The license key file and authentication credentials are provided in an email sent to customer accounts from Cloudera when a new license is issued. If you have an existing license with a CDP Private Cloud Base Edition entitlement, you might not have received an email. In this instance you can identify the authentication credentials from the license key file. If you do not have access to the license key, contact your account representative to receive a copy.

To identify your authentication credentials using your license key file, complete the following steps:

- From cloudera.com, log in to the cloudera.com account associated with the CDP Private Cloud Base license and subscription agreement.
- On the [CDP Private Cloud Base Download page](#), click Download Now and scroll down to the Credential Generator.
- In the Generate Credentials text box, copy and paste the text of the “PGP Signed Message” within your license key file and click Get Credentials. The credentials generator returns your username and password.



Important: Make a note of the authentication credentials. You might need them during installation to complete tasks such as configuring a remote parcel repository, or installing Cloudera Manager packages using a package manager such as YUM, APT, or other tools that you might be using in your environment.

When you obtain your authentication credentials, use them to form the URL where you can access the Cloudera Manager repository in the Cloudera Archive.



Important:

Before performing an upgrade of Cloudera Manager or the CDP Runtime, creating a backup of all the metadata databases is important. This includes the Cloudera Manager database, and the various Runtime component databases such as Hive Metadata Server, Ranger Admin, Ranger KMS, Schema Registry, etc. The backups are necessary if there is a reason to rollback to the prior version.

For information on repositories for the Cloudera Manager 7.11.3 cumulative hotfixes, see [Cumulative hotfixes](#).

The repositories for Cloudera Manager 7.x are listed in the following tables:

Table 2: Cloudera Manager 7.11.3

Repository Type	Repository Location
RHEL 9 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.11.3/redhat9/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.11.3/redhat9/yum/cloudera-manager.repo</pre>
RHEL 8 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.11.3/redhat8/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.11.3/redhat8/yum/cloudera-manager.repo</pre>
RHEL 7 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.11.3/redhat7/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.11.3/redhat7/yum/cloudera-manager.repo</pre>
SLES 15	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.11.3/sles15/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.11.3/sles15/yum/cloudera-manager.repo</pre>
SLES 12	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.11.3/sles12/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.11.3/sles12/yum/cloudera-manager.repo</pre>

Repository Type	Repository Location
Ubuntu 20	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.11.3/ubuntu2004/apt</pre> <p>Repository file:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.11.3/ubuntu2004/apt/cloudera-manager.list</pre>

Table 3: Cloudera Manager 7.10.1

Repository Type	Repository Location
RHEL 8 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.10.1/redhat8/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.10.1/redhat8/yum/cloudera-manager.repo</pre>
RHEL 7 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.10.1/redhat7/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.10.1/redhat7/yum/cloudera-manager.repo</pre>
SLES 12	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.10.1/sles12/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.10.1/sles12/yum/cloudera-manager.repo</pre>
Ubuntu 20	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.10.1/ubuntu2004/apt</pre> <p>Repository file:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.10.1/ubuntu2004/apt/cloudera-manager.list</pre>

Repository Type	Repository Location
Ubuntu 18	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.10.1/ubuntu1804/apt</pre> <p>Repository file:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.10.1/ubuntu1804/apt/cloudera-manager.list</pre>

Table 4: Cloudera Manager 7.9.5

Repository Type	Repository Location
RHEL 8 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.9.5/redhat8/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.9.5/redhat8/yum/cloudera-manager.repo</pre>
RHEL 7 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.9.5/redhat7/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.9.5/redhat7/yum/cloudera-manager.repo</pre>
SLES 12	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.9.5/sles12/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.9.5/sles12/yum/cloudera-manager.repo</pre>
Ubuntu 20	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.9.5/ubuntu2004/apt</pre> <p>Repository file:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.9.5/ubuntu2004/apt/cloudera-manager.list</pre>

Repository Type	Repository Location
Ubuntu 18	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.9.5/ubuntu1804/apt</pre> <p>Repository file:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.9.5/ubuntu1804/apt/cloudera-manager.list</pre>

Table 5: Cloudera Manager 7.8.1

Repository Type	Repository Location
RHEL 8 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.8.1/redhat8/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.8.1/redhat8/yum/cloudera-manager.repo</pre>
RHEL 7 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.8.1/redhat7/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.8.1/redhat7/yum/cloudera-manager.repo</pre>
SLES 12	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.8.1/sles12/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.8.1/sles12/yum/cloudera-manager.repo</pre>
Ubuntu 20	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.8.1/ubuntu2004/apt</pre> <p>Repository file:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.8.1/ubuntu2004/apt/cloudera-manager.list</pre>

Repository Type	Repository Location
Ubuntu 18	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.8.1/ubuntu1804/apt</pre> <p>Repository file:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.8.1/ubuntu1804/apt/cloudera-manager.list</pre>

Table 6: Cloudera Manager 7.7.3-CHF1


Repository Type	Repository Location
RHEL 8 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.7.3-33365545/redhat8/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.7.3-33365545/redhat8/yum/cloudera-manager.repo</pre> <p> Important: You must use Cloudera Manager 7.7.3-CHF1 only when you need to use Python 3.8 for the Cloudera Manager agents. You must install Python 3.8 on all hosts before installing or upgrading to Cloudera Manager 7.7.3-CHF1. Cloudera Manager 7.7.3-CHF1 is only supported with RHEL 7.9, 8.4, and 8.6. For more information, see CDP Private Cloud Base Installation Guide.</p>
RHEL 7.9	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.7.3-33365545/redhat7/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.7.3-33365545/redhat7/yum/cloudera-manager.repo</pre>

Table 7: Cloudera Manager 7.7.3


Repository Type	Repository Location
RHEL 8 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.7.3/redhat8/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.7.3/redhat8/yum/cloudera-manager.repo</pre> <p> Important: You must use Cloudera Manager 7.7.3 only when you need to use Python 3.8 for the Cloudera Manager agents. You must install Python 3.8 on all hosts before installing or upgrading to Cloudera Manager 7.7.3. Cloudera Manager 7.7.3 is only supported with RHEL 8.4 and 8.6. For more information, see CDP Private Cloud Base Installation Guide.</p>

Table 8: Cloudera Manager 7.7.1-CHF3

Repository Type	Repository Location
RHEL 8 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/patch/7.7.1-34818722/redhat8/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/patch/7.7.1-34818722/redhat8/yum/cloudera-manager.repo</pre>
RHEL 7 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/patch/7.7.1-34818722/redhat7/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/patch/7.7.1-34818722/redhat7/yum/cloudera-manager.repo</pre>
SLES 12	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/patch/7.7.1-34818722/sles12/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/patch/7.7.1-34818722/sles12/yum/cloudera-manager.repo</pre>
Ubuntu 20	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/patch/7.7.1-34818722/ubuntu2004/apt</pre> <p>Repository file:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/patch/7.7.1-34818722/ubuntu2004/apt/cloudera-manager.list</pre>

Repository Type	Repository Location
Ubuntu 18	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/patch/7.7.1-34818722/ubuntu1804/apt</pre> <p>Repository file:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/patch/7.7.1-34818722/ubuntu1804/apt/cloudera-manager.list</pre>
IBM PowerPC RHEL 8	<pre>https://username:password@archive.cloudera.com/p/cm7/patch/7.7.1-34818722/redhat8-ppc/yum</pre>
IBM PowerPC RHEL 7	<pre>https://username:password@archive.cloudera.com/p/cm7/patch/7.7.1-34818722/redhat7-ppc/yum</pre>

Table 9: Cloudera Manager 7.7.1

Repository Type	Repository Location
RHEL 8 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.7.1/redhat8/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.7.1/redhat8/yum/cloudera-manager.repo</pre>
RHEL 7 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.7.1/redhat7/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.7.1/redhat7/yum/cloudera-manager.repo</pre>
SLES 12	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.7.1/sles12/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.7.1/sles12/yum/cloudera-manager.repo</pre>

Repository Type	Repository Location
Ubuntu 20	Repository: <pre>https://username:password@archive.cloudera.com/p/cm7/7.7.1/ubuntu2004/apt</pre> Repository file: <pre>https://username:password@archive.cloudera.com/p/cm7/7.7.1/ubuntu2004/apt/cloudera-manager.list</pre>
Ubuntu 18	Repository: <pre>https://username:password@archive.cloudera.com/p/cm7/7.7.1/ubuntu1804/apt</pre> Repository file: <pre>https://username:password@archive.cloudera.com/p/cm7/7.7.1/ubuntu1804/apt/cloudera-manager.list</pre>
IBM PowerPC RHEL 8	<pre>https://username:password@archive.cloudera.com/p/cm7/7.7.1/redhat8-ppc/yum</pre>
IBM PowerPC RHEL 7	<pre>https://username:password@archive.cloudera.com/p/cm7/7.7.1/redhat7-ppc/yum</pre>

Table 10: Cloudera Manager 7.6.5

Repository Type	Repository Location
RHEL 8 Compatible	Repository: <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.5/redhat8/yum</pre> Repository File: <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.5/redhat8/yum/cloudera-manager.repo</pre>
RHEL 7 Compatible	Repository: <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.5/redhat7/yum</pre> Repository File: <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.5/redhat7/yum/cloudera-manager.repo</pre>

Repository Type	Repository Location
SLES 12	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.5/sles12/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.5/sles12/yum/cloudera-manager.repo</pre>
Ubuntu 20	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.5/ubuntu2004/apt</pre> <p>Repository file:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.5/ubuntu2004/apt/cloudera-manager.list</pre>
Ubuntu 18	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.5/ubuntu1804/apt</pre> <p>Repository file:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.5/ubuntu1804/apt/cloudera-manager.list</pre>
IBM PowerPC RHEL 8	<pre>https://username:password@archive.cloudera.com/p/cm7/7.6.5/redhat8-ppc/yum</pre>
IBM PowerPC RHEL 7	<pre>https://username:password@archive.cloudera.com/p/cm7/7.6.5/redhat7-ppc/yum</pre>



Important: Do not upgrade to Cloudera Manager 7.6.1 if you are running CDP Private Cloud Data Services in your deployment.

Table 11: Cloudera Manager 7.6.1 (Service Pack 1 for Cloudera Manager 7.4.4 in CDP Private Cloud Base 7.1.7)

Repository Type	Repository Location
RHEL 8 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.1/redhat8/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.1/redhat8/yum/cloudera-manager.repo</pre>

Repository Type	Repository Location
RHEL 7 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.1/redhat7/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.1/redhat7/yum/cloudera-manager.repo</pre>
SLES 12	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.1/sles12/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.1/sles12/yum/cloudera-manager.repo</pre>
Ubuntu 20	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.1/ubuntu2004/apt</pre> <p>Repository file:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.1/ubuntu2004/apt/cloudera-manager.list</pre>
Ubuntu 18	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.1/ubuntu1804/apt</pre> <p>Repository file:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.1/ubuntu1804/apt/cloudera-manager.list</pre>
IBM PowerPC RHEL 8	<pre>https://username:password@archive.cloudera.com/p/cm7/7.6.1/redhat8-ppc/yum</pre>
IBM PowerPC RHEL 7	<pre>https://username:password@archive.cloudera.com/p/cm7/7.6.1/redhat7-ppc/yum</pre>

Table 12: Cloudera Manager 7.5.5

Repository Type	Repository Location
RHEL 8 Compatible	Repository: <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.5/redhat8/yum</pre> Repository File: <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.5/redhat8/yum/cloudera-manager.repo</pre>
RHEL 7 Compatible	Repository: <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.5/redhat7/yum</pre> Repository File: <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.5/redhat7/yum/cloudera-manager.repo</pre>
SLES 12	Repository: <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.5/sles12/yum</pre> Repository File: <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.5/sles12/yum/cloudera-manager.repo</pre>
Ubuntu 20	Repository: <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.5/ubuntu2004/apt</pre> Repository file: <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.5/ubuntu2004/apt/cloudera-manager.list</pre>
Ubuntu 18	Repository: <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.5/ubuntu1804/apt</pre> Repository file: <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.5/ubuntu1804/apt/cloudera-manager.list</pre>



Note: Cloudera Manager 7.5.4-20668437 contains mitigation for the Apache Log4j vulnerability tracked at CVE-2021-44228. The release mitigates this either by upgrading all dependent libraries to 2.16 log4j or by removing the affected classes. For more information, see [CVE-2021-44228 remediation for CDP Private Cloud Data Services 1.3.3](#)

Table 13: Cloudera Manager 7.5.4-20668437

Repository Type	Repository Location
	Cloudera Manager 7.5.4-20668437 contains mitigation for the Apache Log4j vulnerability tracked at CVE-2021-44228. The release mitigates this either by upgrading all dependent libraries to 2.16 log4j or by removing the affected classes. For more information, see CVE-2021-44228 remediation for CDP Private Cloud Data Services 1.3.3
RHEL 7 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.4-20668437/redhat7/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.4-20668437/redhat7/yum/cloudera-manager.repo</pre>
SLES 12	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.4-20668437/sles12/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.4-20668437/sles12/yum/cloudera-manager.repo</pre>
Ubuntu 18	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.4-20668437/ubuntu1804/apt</pre> <p>Repository file:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.4-20668437/ubuntu1804/apt/cloudera-manager.list</pre>

Table 14: Cloudera Manager 7.5.1

Repository Type	Repository Location
RHEL 7 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.1/redhat7/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.1/redhat7/yum/cloudera-manager.repo</pre>

Repository Type	Repository Location
SLES 12	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.1/sles12/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.1/sles12/yum/cloudera-manager.repo</pre>
Ubuntu 18	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.1/ubuntu1804/apt</pre> <p>Repository file:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.5.1/ubuntu1804/apt/cloudera-manager.list</pre>

Table 15: Cloudera Manager 7.4.4-24429768

Repository Type	Repository Location
	These URLs contain PATCH-5393, which have log4j fixes for Cloudera Manager described in the TSB-545 Critical vulnerability in log4j CVE-2021-44228 .
RHEL 8 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.4.4-24429768/redhat8/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.4.4-24429768/redhat8/yum/cloudera-manager.repo</pre>
RHEL 7 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.4.4-24429768/redhat7/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.4.4-24429768/redhat7/yum/cloudera-manager.repo</pre>
SLES 12	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.4.4-24429768/sles12/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.4.4-24429768/sles12/yum/cloudera-manager.repo</pre>

Repository Type	Repository Location
	These URLs contain PATCH-5393, which have log4j fixes for Cloudera Manager described in the TSB-545 Critical vulnerability in log4j CVE-2021-44228 .
Ubuntu 20	Repository: <pre>https://username:password@archive.cloudera.com/p/cm7/7.4.4-24429768/ubuntu2004/apt</pre> Repository file: <pre>https://username:password@archive.cloudera.com/p/cm7/7.4.4-24429768/ubuntu2004/apt/cloudera-manager.list</pre>
Ubuntu 18	Repository: <pre>https://username:password@archive.cloudera.com/p/cm7/7.4.4-24429768/ubuntu1804/apt</pre> Repository file: <pre>https://username:password@archive.cloudera.com/p/cm7/7.6.1/ubuntu1804/apt/cloudera-manager.list</pre>
IBM PowerPC RHEL 8	<pre>https://username:password@archive.cloudera.com/p/cm7/7.6.1/redhat8-ppc/yum</pre>
IBM PowerPC RHEL 7	<pre>https://username:password@archive.cloudera.com/p/cm7/7.6.1/redhat7-ppc/yum</pre>

Table 16: Cloudera Manager 7.3.1

Repository Type	Repository Location
RHEL 7 Compatible	Repository: <pre>https://username:password@archive.cloudera.com/p/cm7/7.3.1/redhat7/yum</pre> Repository File: <pre>https://username:password@archive.cloudera.com/p/cm7/7.3.1/redhat7/yum/cloudera-manager.repo</pre>
SLES 12	Repository: <pre>https://username:password@archive.cloudera.com/p/cm7/7.3.1/sles12/yum</pre> Repository File: <pre>https://username:password@archive.cloudera.com/p/cm7/7.3.1/sles12/yum/cloudera-manager.repo</pre>

Repository Type	Repository Location
Ubuntu 18	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.3.1/ubuntu1804/apt</pre> <p>Repository file:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.3.1/ubuntu1804/apt/cloudera-manager.list</pre>

Table 17: Cloudera Manager 7.2.4

Repository Type	Repository Location
RHEL 7 Compatible	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.2.4/redhat7/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.2.4/redhat7/yum/cloudera-manager.repo</pre>
SLES 12	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.2.4/sles12/yum</pre> <p>Repository File:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.2.4/sles12/yum/cloudera-manager.repo</pre>
Ubuntu 18	<p>Repository:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.2.4/ubuntu1804/apt</pre> <p>Repository file:</p> <pre>https://username:password@archive.cloudera.com/p/cm7/7.2.4/ubuntu1804/apt/cloudera-manager.list</pre>

Table 18: Cloudera Manager 7.1.4

Repository Type	Repository Location
RHEL 7 Compatible	Repository:
	<code>https://username:password@archive.cloudera.com/p/cm7/7.1.4/redhat7/yum</code>
	Repository file:
	<code>https://username:password@archive.cloudera.com/p/cm7/7.1.4/redhat7/yum/cloudera-manager.repo</code>
SLES 12	Repository:
	<code>https://username:password@archive.cloudera.com/p/cm7/7.1.4/sles12/yum</code>
	Repository file:
	<code>https://username:password@archive.cloudera.com/p/cm7/7.1.4/sles12/yum/cloudera-manager.repo</code>
Ubuntu 18	Repository:
	<code>https://username:password@archive.cloudera.com/p/cm7/7.1.4/ubuntu1804/apt</code>
	Repository file:
	<code>https://username:password@archive.cloudera.com/p/cm7/7.1.4/ubuntu1804/apt/cloudera-manager.list</code>

Table 19: Cloudera Manager 7.1.3

Repository Type	Repository Location
RHEL 7 Compatible	Repository:
	<code>https://username:password@archive.cloudera.com/p/cm7/7.1.3/redhat7/yum</code>
	Repository file:
	<code>https://username:password@archive.cloudera.com/p/cm7/7.1.3/redhat7/yum/cloudera-manager.repo</code>
SLES 12	Repository:
	<code>https://username:password@archive.cloudera.com/p/cm7/7.1.3/sles12/yum</code>
	Repository file:
	<code>https://username:password@archive.cloudera.com/p/cm7/7.1.3/sles12/yum/cloudera-manager.repo</code>

Table 20: Cloudera Manager 7.1.2

Repository Type	Repository Location
RHEL 7 Compatible	Repository:
	<code>https://username:password@archive.cloudera.com/p/cm7/7.1.2/redhat7/yum</code>
	Repository file:
	<code>https://username:password@archive.cloudera.com/p/cm7/7.1.2/redhat7/yum/cloudera-manager.repo</code>

Table 21: Cloudera Manager 7.1.1

Repository Type	Repository Location
RHEL 7 Compatible	Repository:
	<code>https://username:password@archive.cloudera.com/p/cm7/7.1.1/redhat7/yum</code>
	Repository file:
	<code>https://username:password@archive.cloudera.com/p/cm7/7.1.1/redhat7/yum/cloudera-manager.repo</code>

Cloudera Runtime Version Information

Version numbers for current and previous releases of Cloudera Runtime 7.x.

Cloudera Runtime 7.1.9 SP1 is based on Apache Hadoop 3. For more information, see *Cloudera Runtime Component Versions*.

Release date: July 19, 2024

Previous releases:

- Cloudera Runtime 7.1.7 SP3 Release Date: May 07, 2024
- Cloudera Runtime 7.1.9 Release Date: September 08, 2023
- Cloudera Runtime 7.1.7 SP2 Release Date: January 31, 2023
- Cloudera Runtime 7.1.8 Release Date: August 30, 2022
- Cloudera Runtime 7.1.7 SP1 Release Date: March 30, 2022
- Cloudera Runtime 7.1.7 Release Date: August 5, 2021
- Cloudera Runtime 7.1.6 Release Date: March 3, 2021
- Cloudera Runtime 7.1.5 Release Date: November 30, 2020
- Cloudera Runtime 7.1.4 Release Date: October 13, 2020
- Cloudera Runtime 7.1.3 Release Date: August 10, 2020
- Cloudera Runtime 7.1.2 Release Date: July 13, 2020
- Cloudera Runtime 7.1.1 Release Date: May 22, 2020
- Cloudera Runtime 7.0.3 Release Date: November 22, 2019

Related Information

[Cloudera Runtime Component Versions](#)

Cloudera Runtime Download Information

Important: Access to Cloudera Runtime parcels for production purposes requires authentication. To access the parcels at the locations below, you must first have an active subscription agreement and obtain a license key file along with the required authentication credentials (username and password).

The license key file and authentication credentials are provided in an email sent to customer accounts from Cloudera when a new license is issued. If you have an existing license with a CDP Private Cloud Base Edition entitlement, you might not have received an email. In this instance you can identify the authentication credentials from the license key file. If you do not have access to the license key, contact your account representative to receive a copy.

To identify your authentication credentials using your license key file, complete the following steps:

- From cloudera.com, log into the cloudera.com account associated with the CDP Private Cloud Base license and subscription agreement.
- On the [CDP Private Cloud Base Download page](#), click Download Now and scroll down to the Credential Generator.
- In the Generate Credentials text box, copy and paste the text of the “PGP Signed Message” within your license key file and click Get Credentials. The credentials generator returns your username and password.



Important: Make a note of the authentication credentials. You might need them during installation to complete tasks such as configuring a remote parcel repository.



Note: For information on Cloudera Runtime CHF, see the [Cloudera Runtime CHF](#) documentation.

When you obtain your authentication credentials, use them to form the URL where you can access the Runtime repository in the Cloudera Archive. Cloudera Manager can also download the Runtime parcels directly during the installation process.

The repositories for Cloudera Runtime 7.x are listed in the following tables:

Table 22: Cloudera Runtime 7.1.9.1000 (Service Pack 1):

Parcel Repository Location
<code>https://[username]:[password]@archive.cloudera.com/p/cdh7/7.1.9.1000/parcels</code>

Table 23: Cloudera Runtime 7.1.7.3000 (Service Pack 3):

Parcel Repository Location
<code>https://[username]:[password]@archive.cloudera.com/p/cdh7/7.1.7.3000/parcels</code>

Table 24: Cloudera Runtime 7.1.9.0:

Parcel Repository Location
<code>https://[username]:[password]@archive.cloudera.com/p/cdh7/7.1.9.0/parcels</code>

Table 25: Cloudera Runtime 7.1.8.57 (Hotfix for IBM PPC):**Parcel Repository Location**

```
https://[username]:[password]@archive.cloudera.com/p/cdh7/7.1.8.57/parcels
```

Table 26: Cloudera Runtime 7.1.8.10 (Hotfix for IBM PPC):**Parcel Repository Location**

```
https://[username]:[password]@archive.cloudera.com/p/cdh7/7.1.8.10/parcels
```

Table 27: Cloudera Runtime 7.1.8.0:**Parcel Repository Location**

```
https://[username]:[password]@archive.cloudera.com/p/cdh7/7.1.8.0/parcels
```

Table 28: Cloudera Runtime 7.1.7.2000 (Service Pack 2):**Parcel Repository Location**

```
https://[username]:[password]@archive.cloudera.com/p/cdh7/7.1.7.2000/parcels
```

Table 29: Cloudera Runtime 7.1.7.1000 (Service Pack 1):**Parcel Repository Location**

```
https://[username]:[password]@archive.cloudera.com/p/cdh7/7.1.7.1000/parcels
```



Note: Cloudera Runtime 7.1.7 has been replaced by Cloudera Runtime 7.1.7 SP1 for use with CDP Private Cloud Base 7.1.7 SP1.

Table 30: Cloudera Runtime 7.1.7.0 (with log4j fixes in HOTFIX-4836):**Parcel Repository Location**

```
https://[username]:[password]@archive.cloudera.com/p/cdh7/7.1.7.78/parcels
```

Table 31: Cloudera Runtime 7.1.6.0:**Parcel Repository Location**

```
https://[username]:[password]@archive.cloudera.com/p/cdh7/7.1.6.0/parcels
```

Table 32: Cloudera Runtime 7.1.5.0:

Parcel Repository Location
<code>https://[username]:[password]@archive.cloudera.com/p/cdh7/7.1.5.0/parcels</code>

Table 33: Cloudera Runtime 7.1.4.0:

Parcel Repository Location
<code>https://[username]:[password]@archive.cloudera.com/p/cdh7/7.1.4.0/parcels</code>

Table 34: Cloudera Runtime 7.1.3.0:

Parcel Repository Location
<code>https://[username]:[password]@archive.cloudera.com/p/cdh7/7.1.3.0/parcels</code>

Table 35: Cloudera Runtime 7.1.2.1:

Parcel Repository Location
<code>https://[username]:[password]@archive.cloudera.com/p/cdh7/7.1.2.1/parcels</code>

Table 36: Cloudera Runtime 7.1.1.1:

Parcel Repository Location
<code>https://[username]:[password]@archive.cloudera.com/p/cdh7/7.1.1.2001/parcels</code>

Table 37: Cloudera Runtime 7.0.3.0:

Parcel Repository Location
<code>https://[username]:[password]@archive.cloudera.com/p/cdh7/7.0.3.0/parcels</code>

CDP Private Cloud Base Trial Download Information

You can try the CDP Private Cloud Base Edition of Cloudera Data Platform for 60 days without obtaining a license key file.

To download CDP Private Cloud Base without obtaining a license key file, visit the [CDP Private Cloud Base Trial Download](#) page, click Try Now, and follow the download instructions. When you install CDP Private Cloud Base without a license key, you are performing a trial installation that includes an embedded PostgreSQL database and is not suitable for a production environment. For more information on trial installations, see the trial installation documentation.

A 60-day trial of CDP Private Cloud Base Edition can be enabled permanently with the appropriate license. To obtain a CDP Private Cloud Base Edition license, fill in the [Contact Us](#) form or call 866-843-7207

Product Compatibility Matrices

For more information on component compatibility across versions, see the following compatibility matrices:

Replication Manager on CDP Private Cloud Base

CDP Private Cloud Base Replication Manager can replicate HDFS directories, Hive external tables, Impala data, Hive ACID tables, Iceberg tables, Ranger policies and roles for HDFS, Hive, and HBase services, and data in Ozone buckets.

See the following sections for the supported cluster and runtime versions:

- [Replicate from CDH and CDP Private Cloud Base source clusters](#) section lists the cluster and runtime versions to:
 - replicate data from CDH source clusters
 - replicate data between CDP Private Cloud Base clusters using same storage
 - replicate data between CDP Private Cloud Base clusters using different storage
- [Replicate HDFS and Hive data to cloud storage](#)
- [Replicate from HDP 2 and HDP 3 source clusters](#)

Replication policies support the following scenarios:

Kerberos

Replication Manager supports the following replication scenarios when Kerberos authentication is used on a cluster:

- Secure source to a secure destination.
- Insecure source to an insecure destination.
- Insecure source to a secure destination. The following requirements must be met for this scenario:
 - When a destination cluster has multiple source clusters, all the source clusters must either be secure or insecure. Replication Manager does not support a mix of secure and insecure source clusters.
 - The destination cluster must run Cloudera Manager 7.x or higher.
 - The source cluster must run a compatible Cloudera Manager version.
 - This replication scenario requires additional configuration. For more information, see [Replicating from unsecure to secure clusters](#).

Transport Layer Security (TLS)

You can use TLS with Replication Manager. Additionally, Replication Manager supports replication scenarios where TLS is enabled for non-Hadoop services (Hive/Impala) and TLS is disabled Hadoop services (such as HDFS, YARN, and MapReduce).

Apache Knox

When Cloudera Manager is configured with Knox and the source and target clusters are Knox-SSO enabled, you must ensure that you use the Cloudera Manager port in the peer URL when you add the source and target clusters as peers.

Replicate from CDH and CDP Private Cloud Base source clusters

The following tables list the source and destination clusters, lowest supported versions of Cloudera Manager, and the services that are available for each supported cloud provider for CDH and CDP Private Cloud Base source clusters; ensure that the target database name is the same as the source database name, otherwise issues appear during or after data replication:

Table 38: Replicate data from CDH source clusters

Source cluster	Lowest supported source Cloudera Manager version	Lowest supported source Cloudera Runtime version	Lowest supported destination cluster version	Supported services on Replication Manager
CDH 5 CDH 6	6.3.0	5.10	CDP Private Cloud Base 7.0.3	HDFS, Sentry to Ranger*, Hive external tables
*To perform Sentry to Ranger replication using HDFS and Hive external table replication policies, you must have installed Cloudera Manager version 6.3.1 and higher on the source cluster and Cloudera Manager version 7.1.1 and higher on the target cluster.				

Table 39: Replicate data between CDP Private Cloud Base clusters using same storage

Source cluster	Lowest supported source Cloudera Manager version	Lowest supported source Cloudera Runtime version	Destination cluster	Supported services on Replication Manager
CDP Private Cloud Base	7.1.1	7.1.1	CDP Private Cloud Base	<ul style="list-style-type: none"> HDFS Hive external tables
CDP Private Cloud Base	7.7.1	7.1.8	CDP Private Cloud Base	<ul style="list-style-type: none"> Hive ACID tables* Use Cloudera Manager APIs to replicate Ozone buckets.
CDP Private Cloud Base	7.7.1 CHF4	7.1.8	CDP Private Cloud Base	Ozone buckets
CDP Private Cloud Base	7.11.3	7.1.9	CDP Private Cloud Base	<ul style="list-style-type: none"> Iceberg tables Ranger policies and roles, and Ranger audit logs in HDFS**
CDP Private Cloud Base	7.11.3 CHF7	7.1.9 SP1	CDP Private Cloud Base	Atlas replication policies***
<ul style="list-style-type: none"> *You can use REPL commands or Replication Manager to replicate Hive ACID tables between CDP Private Cloud Base 7.1.8 or higher versions using Cloudera Manager versions 7.7.1 or higher. **You can also create Ranger replication policies on Kerberos-enabled CDP Private Cloud Base 7.1.8 or higher clusters using Cloudera Manager 7.7.1 CHF6 and higher, if the Ranger replication feature flag is enabled. ***Replicating Atlas metadata using Hive external table replication policies and Iceberg replication policies, and replicating the metadata and data lineage of all the Hive external tables, Iceberg tables, and any other Atlas supported entities in the source cluster to the target cluster using Atlas replication policies is a technical preview feature. It 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. 				

Table 40: Replicate data between CDP Private Cloud Base clusters using different storage

Source cluster	Lowest supported source Cloudera Manager version	Lowest supported source Cloudera Runtime version	Destination cluster	Supported services on Replication Manager
CDP Private Cloud Base	7.11.3 CHF1	7.1.9	CDP Private Cloud Base	Replicate the data and metadata for Hive external tables from: <ul style="list-style-type: none"> source cluster using HDFS to a target cluster using Dell EMC Isilon storage. source cluster using Dell EMC Isilon storage to a target cluster using HDFS.

Source cluster	Lowest supported source Cloudera Manager version	Lowest supported source Cloudera Runtime version	Destination cluster	Supported services on Replication Manager
CDP Private Cloud Base	7.11.3 CHF2	7.1.9	CDP Private Cloud Base	Replicate Hive ACID tables and Iceberg tables from: <ul style="list-style-type: none"> source cluster using HDFS to a target cluster using Dell EMC Isilon storage. source cluster using Dell EMC Isilon storage to a target cluster using HDFS.
CDP Private Cloud Base	7.11.3 CHF7	7.1.9 SP1	CDP Private Cloud Base	Replicate metadata-only for Ozone storage-backed Hive external tables using Hive external table replication policies. You must replicate the data using Ozone replication policies.



Important: Hive external table replication policies do not support managed to managed table replication. When you replicate from a CDH cluster to a CDP Private Cloud Base cluster, Replication Manager converts managed tables to external tables.

Therefore, to replicate managed tables (ACID) and external tables in a database successfully, you must perform the following steps in the order shown below:

1. Create Hive ACID table replication policy for the database to replicate the managed data.
2. After the replication completes, create the Hive external table replication policy to replicate the external tables in the database.

Ensure that the target cluster version is CDP Private Cloud Base 7.1.8 or higher.

Replicate HDFS and Hive data to cloud storage

CDP Private Cloud Base Replication Manager supports the following replication scenarios:


- Replicate to and from Amazon S3 from CDH 5.14+ and Cloudera Manager version 5.13+.
Replication Manager does not support S3 as a source or destination when S3 is configured to use SSE-KMS.
- Replicate to and from Microsoft ADLS Gen1 from CDH 5.13+ and Cloudera Manager 5.15, 5.16, 6.1+.
- Replicate to Microsoft ADLS Gen2 (ABFS) from CDH 5.13+ and Cloudera Manager 6.1+.
- Supports snapshots from CDH 5.15+ and Cloudera Manager 5.15+.
- Replicate HDFS and Hive external tables from CDP Private Cloud Base 7.1.9 CHF3 and higher clusters using Dell EMC Isilon storage to CDP Public Cloud clusters on AWS, Azure, and GCP.
- Replicate HDFS and Hive external tables from CDP Private Cloud Base 7.1.9 SP1 and higher to CDP Public Cloud clusters on GCP.

Starting in Cloudera Manager 6.1.0, Replication Manager ignores Hive tables backed by Kudu during replication. The change does not affect functionality since Replication Manager does not support tables backed by Kudu. This change was made to guard against data loss due to how the Hive Metastore, Impala, and Kudu interact.

Replicate from HDP 2 and HDP 3 source clusters

Replicating to and from HDP to Cloudera Manager 7.x is not supported by Replication Manager. However, you can replicate data using other methods. The following table lists the methods and the supported data replications to CDP Private Cloud Base clusters that are supported:

Table 41: Replicate data from HDP 2 and HDP 3 source clusters

Lowest supported source version	Services that require alternate replication methods
HDP 2.6.5	HDFS. Use DistCp to replicate data.
HDP 3.1.1	HDFS. Use DistCp to replicate data.
HDP 3.1.1	<ul style="list-style-type: none"> HBase. Use HBase replication to replicate HBase data. Hive external tables. For information to replicate data, contact Cloudera Support.
HDP 3.1.5	Hive ACID tables to CDP 7.1.6 and higher clusters. Use REPL commands to replicate data.  Note: Requires HDP 3.1.5 hotfixes.

Cloudera Manager with Runtime, CDH, and Data Services

Describes which versions of CDH, Cloudera Runtime, and CDP Private Cloud Data Services are supported by Cloudera Manager.



Important:

Before performing an upgrade of Cloudera Manager or the CDP Runtime, creating a backup of all the metadata databases is important. This includes the Cloudera Manager database, and the various Runtime component databases such as Hive Metadata Server, Ranger Admin, Ranger KMS, Schema Registry, etc. The backups are necessary if there is a reason to rollback to the prior version.





Note: Not all combinations of Cloudera Manager, Cloudera Runtime, and CDP Private Cloud Data Services are supported. Ensure that the version of Cloudera Manager you are using supports the version of Cloudera Runtime and CDP Private Cloud Data Services you have selected. For more information, see the [Cloudera Support Matrix](#).

The versions of Cloudera Runtime, CDP Private Cloud Data Services, and CDH clusters that can be managed by Cloudera Manager are limited to the following:

For CDP Private Cloud Base

Table 42: Cloudera Manager support for CDP Private Cloud Base

Cloudera Manager Version	Supported CDH/Runtime versions	Supported CDP Private Cloud Data Services versions
Cloudera Manager 7.11.3 Latest cumulative hotfix	<ul style="list-style-type: none"> Cloudera Runtime 7.1.9 SP1 Cloudera Runtime 7.1.7 SP3 Cloudera Runtime 7.1.9 Cloudera Runtime 7.1.7 SP2 Cloudera Runtime 7.1.8 Cloudera Runtime 7.1.7 SP1 Cloudera Runtime 7.1.7 Cloudera Runtime 7.1.6 Cloudera Runtime 7.1.5 Cloudera Runtime 7.1.4 Cloudera Runtime 7.1.3 Cloudera Runtime 7.1.2 Cloudera Runtime 7.1.1 Cloudera Runtime 7.0.3 CDH 6.3 CDH 6.2 CDH 6.1 CDH 6.0 	None
Cloudera Manager 7.11.3 Cumulative hotfix 5	<ul style="list-style-type: none"> Cloudera Runtime 7.1.7 SP3 Cloudera Runtime 7.1.9 Cloudera Runtime 7.1.7 SP2 Cloudera Runtime 7.1.8 Cloudera Runtime 7.1.7 SP1 Cloudera Runtime 7.1.7 Cloudera Runtime 7.1.6 Cloudera Runtime 7.1.5 Cloudera Runtime 7.1.4 Cloudera Runtime 7.1.3 Cloudera Runtime 7.1.2 Cloudera Runtime 7.1.1 Cloudera Runtime 7.0.3 CDH 6.3 CDH 6.2 CDH 6.1 CDH 6.0 	None
Cloudera Manager 7.11.3  Note: You must install Python 3.8 (or 3.9 for RHEL 9.1) on all hosts before installing or upgrading to Cloudera Manager 7.11.3. For more information, see the Installing Python 3 .	<ul style="list-style-type: none"> Cloudera Runtime 7.1.9 Cloudera Runtime 7.1.7 SP2 Cloudera Runtime 7.1.8 Cloudera Runtime 7.1.7 SP1 Cloudera Runtime 7.1.7 Cloudera Runtime 7.1.6 Cloudera Runtime 7.1.5 Cloudera Runtime 7.1.4 Cloudera Runtime 7.1.3 Cloudera Runtime 7.1.2 Cloudera Runtime 7.1.1 Cloudera Runtime 7.0.3 CDH 6.3 CDH 6.2 CDH 6.1 CDH 6.0 	None




Cloudera Manager Version	Supported CDH/Runtime versions	Supported CDP Private Cloud Data Services versions
<p>Cloudera Manager 7.7.3 should only be used when you need to use Python 3.8 for the Cloudera Manager agents. You must install Python 3.8 on all hosts before installing or upgrading to Cloudera Manager 7.7.3. Cloudera Manager 7.7.3-CHF2 supports only RHEL 8.4, RHEL 8.6, and RHEL 7.9. See the CDP Private Cloud Base Installation Guide for more information.</p>	<ul style="list-style-type: none"> Cloudera Runtime 7.1.8 	None
<p>Cloudera Manager 7.7.1</p> <p> Note: Cloudera recommends you to use latest cumulative hotfix of Cloudera Manager 7.7.1 with Cloudera Runtime 7.1.7-SP2.</p>	<ul style="list-style-type: none"> Cloudera Runtime 7.1.7 SP2 Cloudera Runtime 7.1.8 Cloudera Runtime 7.1.7 SP1 Cloudera Runtime 7.1.7 Cloudera Runtime 7.1.6 Cloudera Runtime 7.1.5 Cloudera Runtime 7.1.4 Cloudera Runtime 7.1.3 Cloudera Runtime 7.1.2 Cloudera Runtime 7.1.1 Cloudera Runtime 7.0.3 CDH 6.3 CDH 6.2 CDH 6.1 CDH 6.0 	None
<p>7.6.7</p> <p> Important: Do not upgrade to Cloudera Manager 7.6.7 if you are running CDP Private Cloud Data Services in your deployment.</p>	<ul style="list-style-type: none"> Cloudera Runtime 7.1.7 SP2 Cloudera Runtime 7.1.7 SP1 Cloudera Runtime 7.1.7 Cloudera Runtime 7.1.6 Cloudera Runtime 7.1.5 Cloudera Runtime 7.1.4 Cloudera Runtime 7.1.3 Cloudera Runtime 7.1.2 Cloudera Runtime 7.1.1 Cloudera Runtime 7.0.3 CDH 6.3 CDH 6.2 CDH 6.1 CDH 6.0 CDH 5.16.2 	None
<p>7.6.1</p> <p> Important: Do not upgrade to Cloudera Manager 7.6.1 if you are running CDP Private Cloud Data Services in your deployment.</p>	<ul style="list-style-type: none"> Cloudera Runtime 7.1.7 SP1 Cloudera Runtime 7.1.7 Cloudera Runtime 7.1.6 Cloudera Runtime 7.1.5 Cloudera Runtime 7.1.4 Cloudera Runtime 7.1.3 Cloudera Runtime 7.1.2 Cloudera Runtime 7.1.1 Cloudera Runtime 7.0.3 CDH 6.3 CDH 6.2 CDH 6.1 CDH 6.0 CDH 5.13 - 5.16 	None







Cloudera Manager Version	Supported CDH/Runtime versions	Supported CDP Private Cloud Data Services versions
7.4.4	<ul style="list-style-type: none"> Cloudera Runtime 7.1.7 Cloudera Runtime 7.1.6 Cloudera Runtime 7.1.5 Cloudera Runtime 7.1.4 Cloudera Runtime 7.1.3 Cloudera Runtime 7.1.2 Cloudera Runtime 7.1.1 Cloudera Runtime 7.0.3 CDH 6.3 CDH 6.2 CDH 6.1 CDH 6.0 CDH 5.13 - 5.16 	None
7.3.1	<ul style="list-style-type: none"> Cloudera Runtime 7.1.6 Cloudera Runtime 7.1.5 Cloudera Runtime 7.1.4 Cloudera Runtime 7.1.3 Cloudera Runtime 7.1.2 Cloudera Runtime 7.1.1 Cloudera Runtime 7.0.3 CDH 6.3 CDH 6.2 CDH 6.1 CDH 6.0 CDH 5.13 - 5.16 	None
7.2.4	<ul style="list-style-type: none"> Cloudera Runtime 7.1.5 Cloudera Runtime 7.1.4 Cloudera Runtime 7.1.3 Cloudera Runtime 7.1.2 Cloudera Runtime 7.1.1 Cloudera Runtime 7.0.3 CDH 6.3 CDH 6.2 CDH 6.1 CDH 6.0 CDH 5.13 - 5.16 	1.2 Supported with Cloudera Runtime 7.1.5 only
7.1.4	<ul style="list-style-type: none"> Cloudera Runtime 7.1.4 Cloudera Runtime 7.1.3 Cloudera Runtime 7.1.2 Cloudera Runtime 7.1.1 Cloudera Runtime 7.0.3 CDH 6.3 CDH 6.2 CDH 6.1 CDH 6.0 CDH 5.13 - 5.16 	None

Cloudera Manager Version	Supported CDH/Runtime versions	Supported CDP Private Cloud Data Services versions
7.1.3	<ul style="list-style-type: none"> Cloudera Runtime 7.1.3 Cloudera Runtime 7.1.2 Cloudera Runtime 7.1.1 Cloudera Runtime 7.0.3 CDH 6.3 CDH 6.2 CDH 6.1 CDH 6.0 CDH 5.13 - 5.16 	1.1
7.1.2	<ul style="list-style-type: none"> Cloudera Runtime 7.1.2 Cloudera Runtime 7.1.1 Cloudera Runtime 7.0.3 CDH 6.3 CDH 6.2 CDH 6.1 CDH 6.0 CDH 5.13 - 5.16 	1.0
7.1.1	<ul style="list-style-type: none"> Cloudera Runtime 7.1.1 Cloudera Runtime 7.0.3 CDH 6.3 CDH 6.2 CDH 6.1 CDH 6.0 CDH 5.13 - 5.16 	None
7.0.3	<ul style="list-style-type: none"> Cloudera Runtime 7.0.3 	None

For CDP Private Cloud Data Services

Table 43: Cloudera Manager support for CDP Private Cloud Data Services

Cloudera Manager Version	Supported CDH/Runtime versions	Supported CDP Private Cloud Data Services versions
7.11.3 cumulative hotfix 4	<ul style="list-style-type: none"> Supported with Cloudera Runtime 7.1.7 SP2, 7.1.8 CHF19 or higher, and 7.1.9 CHF3 only when CDP Private Cloud Data Services is deployed. 	1.5.3  Important: Only supported with Cloudera Runtime 7.1.7 SP2, 7.1.8 CHF19 or higher, and 7.1.9 CHF3.
7.11.3 cumulative hotfix 1	<ul style="list-style-type: none"> Supported with Cloudera Runtime 7.1.7 SP2, 7.1.8 CHF11 or higher, 7.1.9, and 7.1.9 CHF1 only when CDP Private Cloud Data Services is deployed. 	1.5.2  Important: Only supported with Cloudera Runtime 7.1.7 SP2, 7.1.8 CHF11 or higher, 7.1.9, and 7.1.9 CHF1
7.10.1  Important: Upgrade to Cloudera Manager 7.10.1 only if you are running CDP Private Cloud Data Services in your deployment.	<ul style="list-style-type: none"> Supported with Cloudera Runtime 7.1.7 SP2, 7.1.7 SP1, and 7.1.8 CHF4 only when CDP Private Cloud Data Services is deployed. 	1.5.1 Only supported with Cloudera Runtime 7.1.7 SP2, 7.1.7 SP1, and 7.1.8 CHF4.

Cloudera Manager Version	Supported CDH/Runtime versions	Supported CDP Private Cloud Data Services versions
7.9.5  Important: Upgrade to Cloudera Manager 7.9.5 only if you are running CDP Private Cloud Data Services in your deployment.	<ul style="list-style-type: none"> Supported with Cloudera Runtime 7.1.7 SP2, 7.1.7 SP1, and 7.1.8 only when CDP Private Cloud Data Services is deployed. 	1.5.0 Only supported with Cloudera Runtime 7.1.7 SP2, 7.1.7 SP1 and 7.1.8.
7.8.1  Important: Upgrade to Cloudera Manager 7.8.1 only if you are running CDP Private Cloud Data Services in your deployment.	<ul style="list-style-type: none"> Supported with Cloudera Runtime 7.1.7 SP1, and 7.1.8 only when CDP Private Cloud Data Services is deployed. 	1.4.1 Only supported with Cloudera Runtime 7.1.7 SP1 and 7.1.8.
7.6.5  Important: Upgrade to Cloudera Manager 7.6.5 only if you are running CDP Private Cloud Data Services in your deployment.	<ul style="list-style-type: none"> Supported with Cloudera Runtime 7.1.6 , 7.1.7, and 7.1.7 SP1 only when CDP Private Cloud Data Services is deployed. 	1.3.1, 1.3.2, 1.3.3. 1.3.4 are supported with Cloudera Runtime 7.1.6 , 7.1.7. 1.4.0 is supported with Cloudera Runtime 7.1.7 SP1 only.
7.5.5	<ul style="list-style-type: none">  Note: Cloudera Manager 7.5.5 is not compatible with the Spark 3 CDS parcel. Cloudera Runtime 7.1.7 Cloudera Runtime 7.1.6 	1.3.1, 1.3.2, 1.3.3. 1.3.4 Supported with Cloudera Runtime 7.1.6 and 7.1.7 only
7.5.4	<ul style="list-style-type: none">  Note: Cloudera Manager 7.5.4 is not compatible with the Spark 3 CDS parcel. Cloudera Runtime 7.1.7 Cloudera Runtime 7.1.6 	1.3.1, 1.3.2, 1.3.3 Supported with Cloudera Runtime 7.1.6 and 7.1.7 only
7.5.1	<ul style="list-style-type: none">  Note: Cloudera Manager 7.5.1 is not compatible with the Spark 3 CDS parcel. Cloudera Runtime 7.1.7 Cloudera Runtime 7.1.6 	1.3.1 Supported with Cloudera Runtime 7.1.6 and 7.1.7 only

Product Compatibility Matrices for KMS and Encryption Products

Cloudera Navigator encryption comprises several components including Ranger KMS, Navigator Encrypt, KTS, and Key HSM.

The individual compatibility matrices for each component are as follows.

Product Compatibility Matrix for Ranger KMS

Learn about the recommended hardware and supported distributions for Ranger KMS.

Ranger KMS can be installed with either Key Trustee Server or a separate database as the backing keystore.

Ranger KMS: Recommended Hardware and Supported Distributions

Ranger KMS with Key Trustee Server must be installed on a separate host than Key Trustee Server.

The recommended minimum hardware specifications are as follows:

- Processor: 1 GHz 64-bit quad core
- Memory: 8 GB RAM
- Storage: 20 GB on moderate- to high-performance disk drives

Table 44: Ranger KMS Compatibility Matrix

Ranger KMS Version	Supported Operating Systems	Supported Key Trustee Server Versions	Lowest Supported Cloudera Manager Version	Supported CDP Versions
7.1.9	<ul style="list-style-type: none"> RHEL and CentOS: 9.1, 8.8, 8.8 with FIPS, 8.6, 8.4, 8.2, 7.9, 7.8, 7.6, 7.5, 7.4, 7.3, 7.2, 6.10, 6.9, 6.8 Oracle Linux: 7.5, 7.4, 7.3, 7.2, 6.10, 6.9, 6.8 	7.x	7.x	7.x
7.x	<ul style="list-style-type: none"> RHEL and CentOS: 8.6*, 8.4**, 8.2**, 7.9, 7.8, 7.6, 7.5, 7.4, 7.3, 7.2, 6.10, 6.9, 6.8 (* RHEL and CentOS 8.6 is supported only for 7.1.8) (** RHEL and CentOS 8.4, 8.2 are supported only for versions 7.1.7 and higher.) Oracle Linux: 7.5, 7.4, 7.3, 7.2, 6.10, 6.9, 6.8 	7.x	7.x	7.x

Product Compatibility Matrix for Navigator Encrypt

Learn about the recommended hardware and supported distributions for Navigator Encrypt.



Note: Cloudera Enterprise, with the exception of Cloudera Navigator Encrypt, is supported on platforms with Security-Enhanced Linux (SELinux) enabled and in enforcing mode. Cloudera is not responsible for policy support or policy enforcement. If you experience issues with SELinux, contact your OS provider.

Supported command-line interpreters:

- sh (Bourne)
- bash (Bash)
- dash (Ubuntu)



Note: Navigator Encrypt does not support installation or use in chroot environments.

Network Requirements

For new Navigator Key Trustee Server installations, Navigator Encrypt initiates TCP traffic over port 11371 (HTTPS) to the Key Trustee Server.

For new Ranger KMS installations, Navigator Encrypt initiates TCP traffic over port 9494 (HTTPS) to Ranger KMS.

For upgrades, Navigator Encrypt initiates TCP traffic over ports 80 (HTTP) and 443 (HTTPS) to the Navigator Key Trustee Server.

Internet Access

You must have an active connection to the Internet to download many package dependencies, unless you have internal repositories or mirrors containing the dependent packages.

Maintenance Window

Data is not accessible during the encryption process. Plan for system downtime during installation and configuration.

Administrative Access

To enforce a high level of security, all Navigator Encrypt commands require administrative (root) access (including installation and configuration). If you do not have administrative privileges on your server, contact your system administrator before proceeding.

Network Time Protocol (NTP)

The Network Time Protocol (NTP) service synchronizes system time. Cloudera recommends using NTP to ensure that timestamps in system logs, cryptographic signatures, and other auditable events are consistent across systems.

Package Dependencies

Navigator Encrypt requires these packages, which are resolved by your distribution package manager during installation:

- dkms
- keyutils
- openssl
- lsof
- gcc
- cryptsetup

These packages may have other dependencies that are also resolved by your package manager. Installation works with gcc, gcc3, and gcc4.

Table 45: Cloudera Navigator Encrypt Compatibility Matrix

Cloudera Navigator Encrypt Version	Supported Operating Systems	Supported Ranger KMS / Key Trustee Server Versions
7.1.9	<ul style="list-style-type: none"> • RHEL and CentOS: 9.1, 8.8, 8.8 with FIPS, 8.6, 8.4, 8.2, 7.9 • Oracle Linux: 7.7, 7.6 • SLES: 15 SP4 • Ubuntu: 20.04 LTS (Focal) 	Ranger KMS 7.1.9 is supported. KTS 7.1.x is supported.
7.1.8	<ul style="list-style-type: none"> • RHEL and CentOS: 8.6, 8.4, 8.2, 7.9 • Oracle Linux: 7.7, 7.6 • SLES: 12 SP2, SP3, SP4 • Ubuntu: 20.04 LTS (Focal), 18.04 (Bionic) 	7.1.x
7.1.0	<ul style="list-style-type: none"> • RHEL and CentOS: 7.9, 7.8, 7.7, 7.6 • Oracle Linux: 7.7, 7.6 • SLES: 12 SP2, SP3, SP4 • Ubuntu: 16.04 LTS (Xenial), 18 (Bionic) 	7.1.0

Product Compatibility Matrix for KTS and Key HSM

Learn about the recommended hardware and supported distributions for Key Trustee Server and Navigator Key HSM.

Key Trustee Server

Because of a change in the ports used by Key Trustee Server, Navigator Encrypt versions lower than 3.7 and Ranger KMS versions lower than 5.4 are not supported in Key Trustee Server 5.4 and higher.

Recommended Hardware and Supported Distributions

Key Trustee Server must be installed on a dedicated server or virtual machine (VM) that is not used for any other purpose. The backing PostgreSQL database must be installed on the same host as the Key Trustee Server, and must

not be shared with any other services. For high availability, the active and passive Key Trustee Servers must not share physical resources.

The recommended minimum hardware specifications are as follows:

- Processor: 1 GHz 64-bit quad core
- Memory: 8 GB RAM
- Storage: 20 GB on moderate- to high-performance disk drives

Table 46: Cloudera Navigator Key Trustee Server Compatibility Matrix

Cloudera Navigator Key Trustee Server Version	Supported Operating Systems	Lowest Supported Cloudera Manager Version	Lowest Supported Cloudera Navigator Key HSM Versions	Supported Ranger KMS Versions	Supported Cloudera Navigator Encrypt Versions
7.1.9	<ul style="list-style-type: none"> • RHEL and CentOS: 8.8, 8.8 with FIPS, 8.6, 8.4, 8.2, 7.9, 7.8, 7.7, 7.6, 7.5, 7.4, 7.3, 7.2, 6.10, 6.9, 6.8 • Oracle Linux: 7.5, 7.4, 7.3, 7.2, 6.10, 6.9, 6.8 	7.x	7.x	7.x	7.x
7.x	<ul style="list-style-type: none"> • RHEL and CentOS: 8.6*, 8.4**, 8.2**, 7.9, 7.8, 7.7, 7.6, 7.5, 7.4, 7.3, 7.2, 6.10, 6.9, 6.8 (* RHEL and CentOS 8.6 is supported only for 7.1.8) (** RHEL and CentOS 8.4, 8.2 are supported only for versions 7.1.7 and higher.) • Oracle Linux: 7.5, 7.4, 7.3, 7.2, 6.10, 6.9, 6.8 	7.x	7.x	7.x	7.x

Cloudera Navigator Key HSM

Cloudera Navigator Key HSM must be installed on the same host as Key Trustee Server. Although Key HSM is compatible across all versions of Key Trustee Server, Cloudera strongly recommends also upgrading Key HSM after you upgrade Key Trustee Server. See [Installing Cloudera Navigator Key HSM](#) and [Upgrading Cloudera Navigator Key HSM](#) for more information.

Recommended Hardware and Supported Distributions

The following are prerequisites for installing Navigator Key HSM:

- Oracle Java Runtime Environment (JRE) 8 or higher with Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files:

- [JCE for Java SE 8](#)



Note: JDK 1.8u161 and higher enable unlimited strength encryption by default, and do not require policy files.

- OpenJDK 11



Note: The Thales JCE libraries do not support Java 11, so running Key HSM with Thales on OpenJDK 11 is unsupported.

- A supported HSM device:

- Thales (formerly Safenet) Luna

- v6

- HSM firmware version: 6.2.1
 - HSM software version: 5.2.3-1

- v7

- HSM firmware version: 7.0.3
 - HSM software version: 7.2.0

- SafeNet KeySecure

- HSM firmware version: 6.2.1
 - HSM software version: 8.0.1, 8.1.0, 8.7.0

- Thales nSolo, nConnect

- HSM firmware version: 11.4.0
 - Client software version: 2.28.9cam136



Note: Thales Key HSM is unsupported because the Thales client Java libraries do not support Java 11.

- AWS CloudHSM

- Client software version: 1.1.1

- Key Trustee Server 3.8 or higher



Important: You must install Key HSM on the same host as Key Trustee Server.

Root access is required to install Navigator Key HSM.

Table 47: Cloudera Navigator Key HSM Compatibility Matrix

Cloudera Navigator Key HSM Version	Supported Operating Systems	Lowest Supported Key Trustee Server Version
7.x	<ul style="list-style-type: none"> RHEL and CentOS: 8.6*, 8.4**, 8.2**, 7.9, 7.6, 7.5, 7.4, 7.3, 7.2, 6.10, 6.9, 6.8 (* RHEL and CentOS 8.6 is supported only for 7.1.8) (** RHEL and CentOS 8.4, 8.2 are supported only for versions 7.1.7 and higher.) 	7.x

Product Compatibility Matrix for HSM Support

Learn about the HSMs supported for the different CDP/ CDH/ HDP versions.

HSM Support for CDP Private Cloud Base 7.x

Table 48: CDP Private Cloud Base 7.x

	Luna v7	SafeNet KeySecure	nCipher nSolo/nConnect	AWS CloudHSM
Ranger KMS	Luna Client: 7		CipherTrust Manager 2.0 (Supported only for 7.1.8 and 7.1.9)	
Key HSM	HSM firmware version: 7.0.3 HSM software version: 7.2.0	HSM firmware version: 6.2.1 HSM software version: 8.0.1, 8.1.0, 8.7.0	HSM firmware version: 11.4.0 Client software version: 2.28.9cam136 CipherTrust Manager 2.0 (Supported only for 7.1.8 and 7.1.9)	Client software version: 1.1.1

Changes to CDH and HDP Components in CDP Private Cloud Base

CDH and HDP components that have changed or been removed in CDP Private Cloud Base.

Updated CDH Components

CDH Component Changes in CDP Private Cloud Base 7.

Component Changes in CDP Private Cloud Base 7.0

Pig, Flume, Sentry, and Navigator have been removed.

- Pig can be replaced with [Hive](#) or [Spark](#).
- Flume has been replaced with [Cloudera Flow Management](#) (CFM). CFM is a no-code data ingestion and management solution powered by Apache NiFi. Contact your Cloudera account team for more information about moving from Flume to CFM.
- Sentry has been replaced with [Apache Ranger](#). A Sentry-to-Ranger policy migration tool is available for CDP Private Cloud Base 7.1 and migrations will be supported when Replication Manager is used to migrate Hive tables from CDH to CDP.
- Navigator has been replaced with [Apache Atlas](#).
- Cloudera Director is not supported with CDP Private Cloud Base.
- MapReduce v1 was deprecated as of CDH 5.0 and is not supported in CDP.



Note: Replacements are not direct replacements and are only alternate product offerings.

Component Changes in CDP Private Cloud Base 7.1

- The YARN Fair Scheduler is being replaced with the [YARN Capacity Scheduler](#). A migration tool will be provided to convert the Fair Scheduler configurations to Capacity Scheduler.
- The Sentry policy files feature option has been deprecated in CDH 5.x and is entirely removed in CDP Private Cloud Base 7.1. Prior to an upgrade to CDP Private Cloud Base 7.1, the Sentry Service role must be set up and used for enforcing access policies over Hive/Solr. To migrate the existing Sentry Policy file configuration to a Sentry Service, see: [Migrating from Sentry Policy Files to the Sentry Service](#).

- Navigator has been replaced with [Apache Atlas](#).

Navigator lineage data can be transferred to Atlas as part of the CDH to CDP Private Cloud Base upgrade process. Navigator audit data is not transferred to Atlas.

When you upgrade a CDH cluster to CDP Private Cloud Base, the Navigator audits persist. However, services no longer produce audits for Navigator. You can continue to run Navigator to be able to access the audits. See [Transitioning Navigator audits](#).



Note: Replacements are not direct replacements and are only alternate product offerings.

The following CDH components have been included in Cloudera Runtime for the first time in CDP Private Cloud Base 7.1.1.

- [Schema Registry](#)
- [Streams Messaging Manager](#)
- [Streams Replication Manager](#)

Deprecated Items

A deprecated item is a feature, component, platform, or functionality that Cloudera is planning to remove in a future release. Cloudera supports items that are deprecated until they are removed, and the deprecation gives customers time to plan for removal.

The following table lists deprecated items:

Table 49: CDH Components, Subcomponents, and Product Functionality

Item	Related Information	CDH Release in Which Item Is Deprecated	Release in Which Support Is Removed
Apache Crunch	Apache Crunch is deprecated, and will be removed in a future release. Cloudera recommends using Spark 2 instead. Additionally, as of CDH 6.0.0, Crunch is available only as Maven artifacts from the Cloudera Maven repository.	6.0.0	Cloudera Runtime 7
AsyncHBaseSink	AsyncHBaseSink is incompatible with HBase 2.0 and you can no longer use AsyncHBaseSink with Apache Flume. For information about using HBase2Sink with Apache Flume, see .	6.0.0	6.0.0, Cloudera Runtime 7
DataFu		5.9.0	6.0.0, Cloudera Runtime 7
HBaseSink	HBaseSink has been replaced with HBase2Sink, which is compatible with HBase 2.0. For information about using HBase2Sink with Apache Flume.	6.0.0	6.0.0, Cloudera Runtime 7
hbck read-write repair mode	hbck is only available in a read-only inconsistency identifying mode.	6.0.0	6.0.0, Cloudera Runtime 7
HFTP	Use WebHDFS	5.10.1	6.0.0, Cloudera Runtime 7
Hive CLI		5.0.0	Cloudera Runtime 7
HiveServer1		5.3.0	6.0.0, Cloudera Runtime 7
Hue UI version 3	Instead use Hue UI version 4	6.0.0	6.0.0., Cloudera Runtime 7

Item	Related Information	CDH Release in Which Item Is Deprecated	Release in Which Support Is Removed
Key HSM Debug Startup	To get debug information during start up, set the root log level to debug in the /usr/share/keytrustee-server-keyhsm/conf/logback.xml file.	6.1.0	6.1.0, Cloudera Runtime 7
Kite Dataset API	Kite Dataset API is deprecated, and will be removed in a future release.	6.0.0	, Cloudera Runtime 7
Kudu Flume sink configuration parameters	The producer.skipMissingColumn, producer.skipBadColumnValue, and producer.warnUnmatchedRows Kudu Flume sink configuration parameters have been deprecated in favor of producer.missingColumnPolicy, producer.badColumnValuePolicy, and producer.unmatchedRowPolicy respectively.	6.1.0	, Cloudera Runtime 7
kudu perf loadgen tool configuration options	The -table_num_buckets configuration option of the kudu perf loadgen tool is now removed in favor of -table_num_hash_partitions and -table_num_range_partitions	6.1.0	, Cloudera Runtime 7
Legacy Scala clients for Kafka (consumer and producer)	The legacy Scala clients (producer and consumer) that are under the kafka.producer.* and kafka.consumer.* package.	CDK 3.0.0 and CDH 6.0.0	6.1.0, Cloudera Runtime 7
Llama		5.5.0	6.0.0, Cloudera Runtime 7
Mahout		5.5.0	6.0.0
Management of Key Trustee Server without Cloudera Manager		5.9.0	6.0.0, Cloudera Runtime 7
MR Pipes		5.9.0	6.0.0, Cloudera Runtime 7
MRv1, MapReduce v1 APIs, MapReduce service	Migrating from MapReduce 1 to MapReduce 2	5.0.0	6.0.0, Cloudera Runtime 7
Navigator Encrypt Filesystem-Level Encryption Using eCryptfs	There is no support for creating new eCryptfs mount points. Previously existing eCryptfs mount points are not affected.	July 2015	6.0.0, Cloudera Runtime 7
Navigator Encrypt migration command	The navencrypt-migration command is deprecated, and has been removed.	February 1, 2018	6.0.0, Cloudera Runtime 7
Old NameNode UI		5.5.0	6.0.0, Cloudera Runtime 7
Oozie Hive Action	Oozie Hive 2 Action Extension	5.7.0	6.0.0, Cloudera Runtime 7
Parquet library with group ID com.twitter		6.0.0	6.0.0, Cloudera Runtime 7
Parquet methods for reading metadata on the client side	See API Methods Removed .	6.0.0	6.0.0, Cloudera Runtime 7
Python 2.6	Python 2.6, packaged with RHEL6, is deprecated. Key Trustee Server uses the utilities ktadmin and keytrustee-orgtool, which use the native version of Python that is packaged with the host OS.	6.0.0	To be determined.

Item	Related Information	CDH Release in Which Item Is Deprecated	Release in Which Support Is Removed
Sentry policy files	Migrating from Sentry Policy Files to the Sentry Service	5.8.0	6.0.0, Cloudera Runtime 7
Spark 1.x		5.13	6.0.0, Cloudera Runtime 7
Spark Standalone		5.5.0	6.0.0, Cloudera Runtime 7
Sqoop2		5.9.0	6.0.0, Cloudera Runtime 7
Unmanaged (CLI-based) CDH deployments		6.0.0	6.0.0, Cloudera Runtime 7
Whirr		5.5.0	6.0.0, Cloudera Runtime 7
YARN Fair Scheduler	See Migrating from Fair Scheduler to Capacity Scheduler	5.9.0	6.0.0, Cloudera Runtime 7

Updated HDP Components

HDP Component Changes in CDP Private Cloud Base 7.

Component Changes in CDP Private Cloud Base 7.0

Flume, Storm, Druid, Falcon, Mahout, and Ambari have been removed.

- Flume workloads can be migrated to [Cloudera Flow Management](#) (CFM). CFM is a no-code data ingestion and management solution powered by Apache NiFi. Contact your Cloudera account team for more information about moving from Flume to CFM.
- Storm can be replaced with [Cloudera Streaming Analytics](#) (CSA) powered by Apache Flink. Contact your Cloudera account team for more information about moving from Storm to CSA.
- [Cloudera Manager](#) has replaced Ambari, but Ambari will be used as part of the upgrade process from HDP to CDP Private Cloud Base 7.1.



Note: Replacements are not direct replacements and are only alternate product offerings.

Component Changes in CDP Private Cloud Base 7.1

Data Lifecycle Manager (DLM), Data Steward Studio (DSS), Hive LLAP, and Streaming Analytics Manager (SAM) are being removed.

- Data Lifecycle Manager (DLM) is being replaced with [Replication Manager](#).
- Data Steward Studio (DSS) is being replaced with [Data Catalog](#).
- Hive LLAP is being replaced with [Cloudera Data Warehouse](#) (CDW) as CDW is not yet available on CDP Private Cloud Base.
- Streaming Analytics Manager (SAM) was closely tied with Storm and has been removed.
- Cloudbreak is not supported in CDP Private Cloud Base and must not be present when upgrading.
- DataPlane is not supported in CDP Private Cloud Base and must not be present when upgrading.



Note: Replacements are not direct replacements and are only alternate product offerings.

The following HDF components were previously available for installation on HDP, and - have been included in Cloudera Runtime for the first time in CDP Private Cloud Base 7.1.1.

- [Schema Registry](#)
- [Streams Messaging Manager](#)
- [Streams Replication Manager](#)

HDP Core component version changes

You must be aware of the version number changes for the core components included in HDP 2.6.5.x.

Component	HDP 2.6.5.x	Ambari 7.1.x.x	CDP Private Cloud Base (CM)
Hadoop	2.7.3	3.1.1	3.1.1
Hive	1.2.1 2.1.0 (LLAP)	3.1.3000 CLI Removed	3.1.3000
Hive LLAP	available	removed	removed
Hive MR	1.2.1	removed	removed
Hive TEZ	0.7.0	0.9.1	0.9.1
Oozie	4.2.0	5.1.0	5.1.0
Sqoop	1.4.6	1.4.7	1.4.7
Flume	1.5.2	removed	removed
Pig	0.16.0	removed	removed
HBase	1.1.2	2.2.3	2.2.3
Phoenix	4.7.0	5.0.0	5.0.0
Knox	0.12.0	1.3.0	1.3.0
Ranger	0.7.0	2.0.0	2.0.0
Ranger KMS	0.7.0	2.0.0	2.0.0
Druid	0.10.1	Not available	Not available
Storm	1.1.0	removed	removed
Spark	1.6.3 2.3.0	2.4.5	2.4.5
ZooKeeper	3.4.6	3.5.5	3.5.5
Zeppelin	0.7.3	0.8.2	0.8.2
Falcon	0.10.0	removed	removed
Atlas	0.8.0	2.0.0	2.0.0
Kafka	1.0.0	2.4.0	2.4.0

Some of the components are removed between HDP and Ambari PvC Base. You must work with the application teams to transition workloads to another part of the stack before you upgrade. When the transition is complete, remove those components from the HDP stack before you upgrade to Ambari PvC Base.

Changes to Ambari and HDP services

Review the list of additional components that are added to your cluster, along with the deprecated services and views that are removed during the process of upgrading to Ambari 7.1.x and HDP intermediate bits.

These services are removed automatically as part of the upgrade process. You can also remove these services manually prior to the upgrade by following the steps below:

1. Log in to the Ambari UI.
2. Click the service you wish to remove.
3. From the Service Actions menu, click Delete Service.

Ambari 2.6.2.x to Ambari 7.1.x.x.

The Ambari 2.6.2.x to Ambari 7.1.x.x upgrade removes the following views:

- Hive View 1.5, Hive View 2
- Hue To Ambari View Migration
- Slider
- Storm
- Tez
- Pig

HDP 2.6.5.x to HDP intermediate bits

The HDP 2.6.5.x to HDP 7.1.x.x upgrade adds the following components if YARN is deployed in the cluster being upgraded:

- YARN Registry DNS
- YARN Timeline Service V2.0 Reader

The HDP 2.6.5.x to HDP intermediate bits upgrade removes the following services:

- Druid
- Superset
- Accumulo
- Flume
- Mahout
- Falcon
- Spark 1.6
- Slider
- WebHCat
- Spark Thrift Server



Caution:

- Druid is not supported on Ambari 7.1.1. Druid is supported from Ambari 7.1.2 version onwards. You can upgrade when Druid is available on CDP Private Cloud Base.
- When you upgrade Ambari to a later version, Accumulo service is not supported as Accumulo is not compatible with other components and cannot access the data. Hence, you must delete Accumulo service when upgrading Ambari.



Note: You must remove Accumulo on the Ambari managed HDP cluster as the AM2CM tool does not transition the Accumulo component from HDP to CDP. However, you can add Accumulo as a service on the Cloudera manager managed CDP Private Cloud Base cluster.



Important:

- Upgrading from HDP 2.6.5.x to HDP intermediate bits does not support HDP Search to Cloudera Search upgrade.
- LLAP is not supported in CDP Private Cloud Base. You can move the LLAP workloads to CDW Public cloud or Private Cloud.
- The Hive View and the Tez View (collectively known as Ambari Views) are deprecated in CDP Private Cloud Base.

Assessing the Impact of Apache Hive

Apache Hive key features

Major changes to Apache Hive 2.x improve Apache Hive 3.x transactions and security. Knowing the major differences between these versions is critical for SQL users, including those who use Apache Spark and Apache Impala.

Hive is a data warehouse system for summarizing, querying, and analyzing huge, disparate data sets. Cloudera Runtime (CR) services include Hive on Tez and Hive Metastore. Hive on Tez is based on Apache Hive 3.x, a SQL-based data warehouse system. The enhancements in Hive 3.x over previous versions can improve SQL query performance, security, and auditing capabilities. The Hive metastore (HMS) is a separate service, not part of Hive, not even necessarily on the same cluster. HMS stores the metadata on the backend for Hive, Impala, Spark, and other components.

ACID transaction processing

Hive 3 tables are ACID (Atomicity, Consistency, Isolation, and Durability)-compliant. Hive 3 write and read operations improve the performance of transactional tables. Atomic operations include simple writes and inserts, writes to multiple partitions, and multiple inserts in a single SELECT statement. A read operation is not affected by changes that occur during the operation. You can insert or delete data, and it remains consistent throughout software and hardware crashes. Creation and maintenance of Hive tables is simplified because there is no longer any need to bucket tables.

Shared Hive metastore

Cloudera Runtime (CR) services include Hive and Hive Metastore (HMS). HMS supports the interoperability of multiple compute engines, Impala and Spark for example. HMS simplifies access between various engines and user data access.

Scheduled Queries

Using SQL statements, you can schedule Hive queries to run on a recurring basis, monitor query progress, temporarily ignore a query schedule, and limit the number running in parallel. You can use scheduled queries to start compaction and periodically rebuild materialized views, for example.

Low-latency analytical processing (CDP Public Cloud)

Hive processes transactions using low-latency analytical processing (LLAP) or the Apache Tez execution engine. The Hive LLAP service is not available in CDP Private Cloud Base.

Spark integration with Hive

Spark and Hive tables interoperate using the Hive Warehouse Connector and Spark Direct Reader to access ACID managed tables. You can access external tables from Spark directly using SparkSQL.

You do not need HWC to read or write Hive external tables. Spark users just read from or write to Hive directly. You can read Hive external tables in ORC or Parquet formats. You can write Hive external tables in ORC format only. (See link below.)

Security improvements

Apache Ranger secures Hive data by default. To meet demands for concurrency improvements, ACID support, render security, and other features, Hive tightly controls the location of the warehouse on a file system, or object store, and memory resources.

With Apache Ranger and Apache Hive ACID support, your organization will be ready to support and implement GDPR (General Data Protection Regulation).

Workload management at the query level

You can configure who uses query resources, how much can be used, and how fast Hive responds to resource requests. Workload management can improve parallel query execution, cluster sharing for queries, and query performance. Although the names are similar, Hive workload management queries are unrelated to Cloudera Workload XM for reporting and viewing millions of queries and hundreds of databases.

Materialized views

Because multiple queries frequently need the same intermediate roll up or joined table, you can avoid costly, repetitious query portion sharing, by precomputing and caching intermediate tables into views.

Query results cache

Hive filters and caches similar or identical queries. Hive does not recompute the data that has not changed. Caching repetitive queries can reduce the load substantially when hundreds or thousands of users of BI tools and web services query Hive.

Connection Pooling

Hive supports HikariCP JDBC connection pooling.

Apache Hive 3 architectural overview

Understanding Apache Hive 3 major design features, such as default ACID transaction processing, can help you use Hive to address the growing needs of enterprise data warehouse systems.

Apache Tez

Apache Tez is the Hive execution engine for the Hive on Tez service, which includes HiveServer (HS2) in Cloudera Manager. MapReduce is not supported. In a Cloudera cluster, if a legacy script or application specifies MapReduce for execution, an exception occurs. Most user-defined functions (UDFs) require no change to execute on Tez instead of MapReduce.

With expressions of directed acyclic graphs (DAGs) and data transfer primitives, execution of Hive queries on Tez instead of MapReduce improves query performance. In Cloudera Data Platform (CDP), Tez is usually used only by Hive, and launches and manages Tez AM automatically when Hive on Tez starts. SQL queries you submit to Hive are executed as follows:

- Hive compiles the query.
- Tez executes the query.
- Resources are allocated for applications across the cluster.
- Hive updates the data in the data source and returns query results.

Hive on Tez runs tasks on ephemeral containers and uses the standard YARN shuffle service.

Data storage and access control

One of the major architectural changes to support Hive 3 design gives Hive much more control over metadata memory resources and the file system, or object store. The following architectural changes from Hive 2 to Hive 3 provide improved security:

- Tightly controlled file system and computer memory resources, replacing flexible boundaries: Definitive boundaries increase predictability. Greater file system control improves security.
- Optimized workloads in shared files and YARN containers

CDP Private Cloud Base stores Hive data on HDFS by default. CDP Public Cloud stores Hive data on S3 by default. In the cloud, Hive uses HDFS merely for storing temporary files. Hive 3 is optimized for object stores such as S3 in the following ways:

- Hive uses ACID to determine which files to read rather than relying on the storage system.
- In Hive 3, file movement is reduced from that in Hive 2.
- Hive caches metadata and data aggressively to reduce file system operations

The major authorization model for Hive is Ranger. Hive enforces access controls specified in Ranger. This model offers stronger security than other security schemes and more flexibility in managing policies.

This model permits only Hive to access the data warehouse. If you do not enable the Ranger security service, or other security, CDP Private Cloud Base by default Hive uses storage-based authorization (SBA) based on user impersonation.

HDFS permission changes

In CDP Private Cloud Base, SBA relies heavily on HDFS access control lists (ACLs). ACLs are an extension to the permissions system in HDFS. CDP Private Cloud Base turns on ACLs in HDFS by default, providing you with the following advantages:

- Increased flexibility when giving multiple groups and users specific permissions
- Convenient application of permissions to a directory tree rather than by individual files

Transaction processing

You can deploy new Hive application types by taking advantage of the following transaction processing characteristics:

- Mature versions of ACID transaction processing:

ACID tables are the default table type.

ACID enabled by default causes no performance or operational overload.

- Simplified application development, operations with strong transactional guarantees, and simple semantics for SQL commands

You do not need to bucket ACID tables.

- Materialized view rewrites
- Automatic query cache
- Advanced optimizations

Hive client changes

CDP Private Cloud Base supports the thin client Beeline for working on the command line. You can run Hive administrative commands from the command line. Beeline uses a JDBC connection to Hive on Tez to execute commands. Parsing, compiling, and executing operations occur in Hive on Tez. Beeline supports many of the command-line options that Hive CLI supported. Beeline does not support `hive -e set key=value` to configure the Hive Metastore.

You enter supported Hive CLI commands by invoking Beeline using the `hive` keyword, command option, and command. For example, `hive -e set`. Using Beeline instead of the thick client Hive CLI, which is no longer supported, has several advantages, including low overhead. Beeline does not use the entire Hive code base. A small number of daemons required to run queries simplifies monitoring and debugging.

Hive enforces allowlist and denylist settings that you can change using SET commands. Using the denylist, you can restrict memory configuration changes to prevent instability. Different Hive instances with different allowlists and denylists to establish different levels of stability.

Apache Hive Metastore sharing

Hive, Impala, and other components can share a remote Hive metastore. In CDP Public Cloud, HMS uses a pre-installed MySQL database. You perform little, or no, configuration of HMS in the cloud.

Spark integration

Spark and Hive tables interoperate using the Hive Warehouse Connector.

You can access ACID and external tables from Spark using the Hive Warehouse Connector. You do not need the Hive Warehouse Connector to read Hive external tables from Spark and write Hive external tables from Spark. You do not need HWC to read or write Hive external tables. Spark users just read from or write to Hive directly. You can read Hive external tables in ORC or Parquet formats. You can write Hive external tables in ORC format only. (See [link below](#).)

Query execution of batch and interactive workloads

You can connect to Hive using a JDBC command-line tool, such as Beeline, or using an JDBC/ODBC driver with a BI tool, such as Tableau. Clients communicate with an instance of the same Hive on Tez version. You configure the settings file for each instance to perform either batch or interactive processing.

Key semantic changes and workarounds

As SQL Developer, Analyst, or other Hive user, you need to know potential problems with queries due to semantic changes. Some of the operations that changed were not widely used, so you might not encounter any of the problems associated with the changes.

Over the years, Apache Hive committers enhanced versions of Hive supported in legacy releases of CDH and HDP, with users in mind. Changes were designed to maintain compatibility with Hive applications. Consequently, few syntax changes occurred over the years. A number of semantic changes, described in this section did occur, however. Workarounds are described for these semantic changes.

Casting timestamps

Results of applications that cast numerics to timestamps differ from Hive 2 to Hive 3. Apache Hive changed the behavior of CAST to comply with the SQL Standard, which does not associate a time zone with the TIMESTAMP type.

Before Upgrade to CDP

Casting a numeric type value into a timestamp could be used to produce a result that reflected the time zone of the cluster. For example, 1597217764557 is 2020-08-12 00:36:04 PDT. Running the following query casts the numeric to a timestamp in PDT:

```
> SELECT CAST(1597217764557 AS TIMESTAMP);
| 2020-08-12 00:36:04 |
```

After Upgrade to CDP

Casting a numeric type value into a timestamp produces a result that reflects the UTC instead of the time zone of the cluster. Running the following query casts the numeric to a timestamp in UTC.

```
> SELECT CAST(1597217764557 AS TIMESTAMP);
| 2020-08-12 07:36:04.557 |
```

Action Required

Change applications. Do not cast from a numeral to obtain a local time zone. Built-in functions `from_utc_timestamp` and `to_utc_timestamp` can be used to mimic behavior before the upgrade.

Casting invalid dates

Casting of an invalid date differs from Hive 1 in CDH 5 to Hive 3 in CDP. Hive 3 uses a different parser formatter from the one used in Hive 1, which affects semantics. Hive 1 considers 00 invalid for date fields. Hive 3 considers 00 valid for date fields. Neither Hive 1 nor Hive 3 correctly handles invalid dates, and Hive-25056 addresses this issue.

Before Upgrade to CDP

Casting of invalid date (zero value in one or more of the 3 fields of date, month, year) returns a NULL value:

```
SELECT CAST ( '0000-00-00' as date) , CAST ( '000-00-00 00:00:00' AS TIMESTAMP ) ;
```

After Upgrade to CDP

Casting of an invalid date returns a result.

```
> SELECT CAST ( '0000-00-00' as date) , CAST ( '000-00-00 00:00:00' AS TIMESTAM
MP) ;
...
00002-11-30 00:00:00.0
```

Action Required

Do not cast invalid dates in Hive 3.

Changing incompatible column types

A default configuration change can cause applications that change column types to fail.

Before Upgrade to CDP

In HDP 2.x and CDH 5.x and CDH 6 hive.metastore.disallow.incompatible.col.type.changes is false by default to allow changes to incompatible column types. For example, you can change a STRING column to a column of an incompatible type, such as MAP<STRING, STRING>. No error occurs.

After Upgrade to CDP

In CDP, hive.metastore.disallow.incompatible.col.type.changes is true by default. Hive prevents changes to incompatible column types. Compatible column type changes, such as INT, STRING, BIGINT, are not blocked.

Action Required

Change applications to disallow incompatible column type changes to prevent possible data corruption. Check ALTER TABLE statements and change those that would fail due to incompatible column types.

Understanding CREATE TABLE behavior

Hive table creation has changed significantly since Hive 3 to improve useability and functionality. If you are upgrading from CDH or HDP, you must understand the changes affecting legacy table creation behavior.

Hive has changed table creation in the following ways:

- Creates ACID-compliant table, which is the default in CDP
- Supports simple writes and inserts
- Writes to multiple partitions
- Inserts multiple data updates in a single SELECT statement
- Eliminates the need for bucketing.

If you have an ETL pipeline that creates tables in Hive, the tables will be created as ACID. Hive now tightly controls access and performs compaction periodically on the tables. Using ACID-compliant, transactional tables causes no performance or operational overload. The way you access managed Hive tables from Spark and other clients changes. In CDP, access to external tables requires you to set up security access permissions.

You must understand the behavior of the CREATE TABLE statement in legacy platforms like CDH or HDP and how the behavior changes after you upgrade to CDP.

Before upgrading to CDP

In CDH 5, CDH 6, and HDP 2, by default CREATE TABLE creates a non-ACID managed table in plain text format.

In HDP 3 and CDP 7.1.0 through 7.1.7.x, by default CREATE TABLE creates either a full ACID transactional table in ORC format or insert-only ACID transactional tables for all other table formats.

After upgrading to CDP

- If you are upgrading from HDP 2, CDH 5, or CDH 6 to CDP 7.1.0 through CDP 7.1.8, by default CREATE TABLE creates a full ACID transactional table in ORC format or insert-only ACID transactional tables for all other table formats.
- If you are upgrading from HDP 3 or CDP 7.1.0 through 7.1.7.x to CDP 7.1.8, the existing behavior persists and CREATE TABLE creates either a full ACID transactional table in ORC format or insert-only ACID transactional tables for all other table formats.

Now that you understand the behavior of the CREATE TABLE statement, you can choose to modify the default table behavior by configuring certain properties. The order of preference for configuration is as follows:

Override default behavior when creating the table

Irrespective of the database, session, or site-level settings, you can override the default table behavior by using the MANAGED or EXTERNAL keyword in the CREATE TABLE statement.

```
CREATE [MANAGED][EXTERNAL] TABLE foo (id INT);
```

Set the default table type at a database level

You can use the database property, defaultTableType=EXTERNAL or ACID to specify the default table type to be created using the CREATE TABLE statement. You can specify this property when creating the database or at a later point using the ALTER DATABASE statement. For example:

```
CREATE DATABASE test_db WITH DBPROPERTIES ('defaultTableType'='EXTERNAL');
```

In this example, tables created under the test_db database using the CREATE TABLE statement creates external tables with the purge functionality enabled (external.table.purge = 'true').

You can also choose to configure a database to allow only external tables to be created and prevent creation of ACID tables. While creating a database, you can set the database property, EXTERNAL_TABLES_ONLY=true to ensure that only external tables are created in the database. For example:

```
CREATE DATABASE test_db WITH DBPROPERTIES ('EXTERNAL_TABLES_ONLY'='true');
```

Set the default table type at a session level

You can configure the CREATE TABLE behavior within an existing beeline session by setting hive.create.as.external.legacy to true or false. Setting the value to true results in configuring the CREATE TABLE statement to create external tables by default.

When the session ends, the default CREATE TABLE behavior also ends.

Set the default table type at a site level

You can configure the CREATE TABLE behavior at the site level by configuring the hive.create.as.insert.only and hive.create.as.acid properties in Cloudera Manager. When configured at the site level, the behavior persists from session to session. For more information, see Configuring CREATE TABLE behavior.

If you are a Spark user, switching to legacy behavior is unnecessary. Calling 'create table' from SparkSQL, for example, creates an external table after upgrading to CDP as it did before the upgrade. You can connect to Hive using the Hive Warehouse Connector (HWC) to read Hive ACID tables from Spark. To write ACID tables to Hive from Spark, you use the HWC and HWC API. Spark creates an external table with the purge property when you do not use the HWC API. For more information, see Hive Warehouse Connector for accessing Spark data.

Configuring legacy CREATE TABLE behavior

After you upgrade to CDP Private Cloud Base and migrate old tables, the legacy CREATE TABLE behavior of Hive is no longer available by default and you might want to switch to the legacy behavior. Legacy behavior might solve compatibility problems with your scripts during data migration, for example, when running ETL.

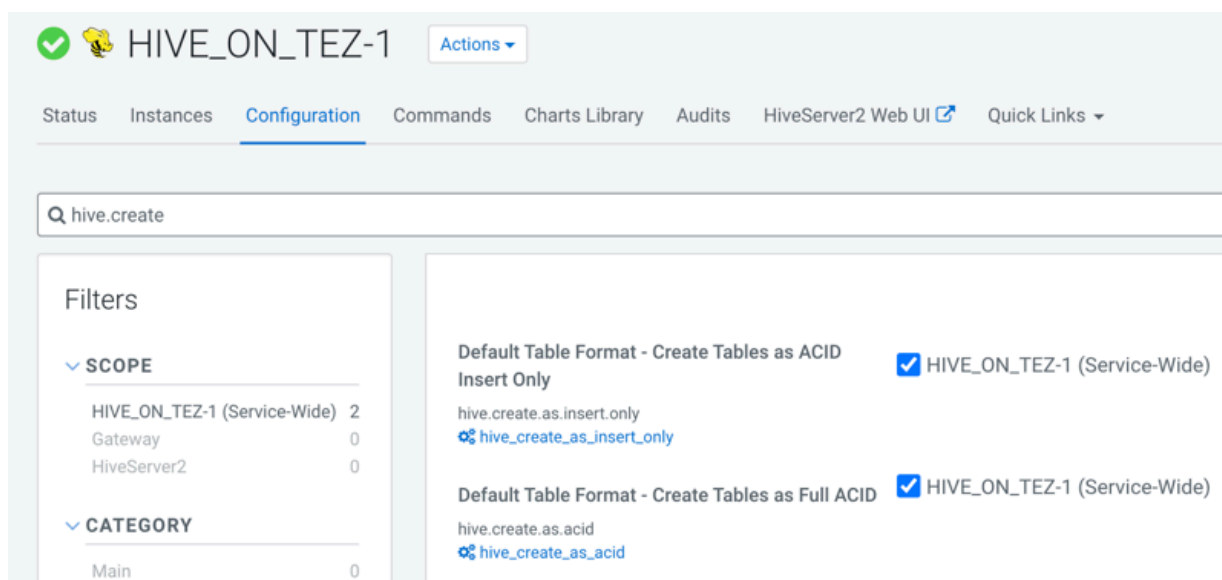
About this task

In CDP, running a CREATE TABLE statement by default creates a full ACID table for ORC file format and insert-only ACID table for other file formats. You can change the default behavior to use the legacy CREATE TABLE behavior. When you configure legacy behavior, CREATE TABLE creates external tables with the purge functionality enabled (`external.table.purge = 'true'`). Therefore, when the table is dropped, data is also deleted from the file system.

You can configure legacy CREATE TABLE behavior at the site level by configuring properties in Cloudera Manager. When configured at the site level, the behavior persists from session to session.

Procedure

1. In Cloudera Manager, click Clusters and select the Hive on Tez service.
2. From the Hive on Tez service, go to the Configuration tab and search for `hive.create`.



3. If the following properties are selected, clear the selection to enable legacy CREATE TABLE behavior.
 - Default Table Format - Create Tables as ACID Insert Only (`hive.create.as.insert.only`)
 - Default Table Format - Create Tables as Full ACID (`hive.create.as.acid`)

Results

Legacy behavior is enabled and the CREATE TABLE statement now creates external tables with the `external.table.purge` property set to true.

Handling table reference syntax

For ANSI SQL compliance, Hive 3.x rejects ``db.table`` in SQL queries as described by the Hive-16907 bug fix. A dot (.) is not allowed in table names. As a Data Engineer, you need to ensure that Hive tables do not contain these references before migrating the tables to CDP, that scripts are changed to comply with the SQL standard references, and that users are aware of the requirement.

About this task

To change queries that use such ``db.table`` references thereby preventing Hive from interpreting the entire `db.table` string incorrectly as the table name, you enclose the database name and the table name in backticks as follows:

A dot (.) is not allowed in table names.

Procedure

1. Find a table having the problematic table reference.
For example, `math.students` appears in a `CREATE TABLE` statement.
2. Enclose the database name and the table name in backticks.

```
CREATE TABLE `math`.`students` (name VARCHAR(64), age INT, gpa DECIMAL(3,2));
```

Add Backticks to Table References

CDP includes the Hive-16907 bug fix, which rejects ``db.table`` in SQL queries. A dot (.) is not allowed in table names. You need to change queries that use such references to prevent Hive from interpreting the entire `db.table` string as the table name.

Procedure

1. Find a table having the problematic table reference.

```
math.students
```

appears in a `CREATE TABLE` statement.

2. Enclose the database name and the table name in backticks.

```
CREATE TABLE `math`.`students` (name VARCHAR(64), age INT, gpa DECIMAL(3,2));
```

Handling the Keyword APPLICATION

If you use the keyword `APPLICATION` in your queries, you might need to modify the queries to prevent failure.

To prevent a query that uses a keyword from failing, enclose the query in backticks.

Before Upgrade to CDP

In CDH releases, such as CDH 5.13, queries that use the word `APPLICATION` in queries execute successfully. For example, you could use this word as a table name.

```
> select f1, f2 from application
```

After Upgrade to CDP

A query that uses the keyword `APPLICATION` fails.

Action Required

Change applications. Enclose queries in backticks. `SELECT field1, field2 FROM `application`;`

Dropping partitions

The `OFFLINE` and `NO_DROP` keywords in the `CASCADE` clause for dropping partitions causes performance problems and is no longer supported.

Before Upgrade to CDP

You could use `OFFLINE` and `NO_DROP` keywords in the `DROP CASCADE` clause to prevent partitions from being read or dropped.

After Upgrade to CDP

OFFLINE and NO_DROP are not supported in the DROP CASCADE clause.

Action Required

Change applications to remove OFFLINE and NO_DROP from the DROP CASCADE clause. Use an authorization scheme, such as Ranger, to prevent partitions from being dropped or read.

Handling output of greatest and least functions

To calculate the greatest (or least) value in a column, you need to work around a problem that occurs when the column has a NULL value.

Before Upgrade to CDP

The greatest function returned the highest value of the list of values. The least function returned the lowest value of the list of values.

After Upgrade to CDP

Returns NULL when one or more arguments are NULL.

Action Required

Use NULL filters or the nvl function on the columns you use as arguments to the greatest or least functions.

```
SELECT greatest(nvl(col1,default value incase of NULL),nvl(col2,default value incase of NULL));
```

Renaming tables

To harden the system, Hive data can be stored in HDFS encryption zones. RENAME has been changed to prevent moving a table outside the same encryption zone or into a no-encryption zone.

Before Upgrade to CDP

In CDH and HDP, renaming a managed table moves its HDFS location.

After Upgrade to CDP

Renaming a managed table moves its location only if the table is created without a LOCATION clause and is under its database directory.

Action Required

None

TRUNCATE TABLE on an external table

Hive 3 does not support TRUNCATE TABLE on external tables. Truncating an external table results in an error. You can truncate an external table if you change your applications to set a table property to purge data.

Before Upgrade to CDP

Some legacy versions of Hive supported TRUNCATE TABLE on external tables.

After Upgrade to CDP Private Cloud Base

By default, TRUNCATE TABLE is supported only on managed tables. Attempting to truncate an external table results in the following error:

```
Error: org.apache.spark.sql.AnalysisException: Operation not allowed: TRUNCATE TABLE on external tables
```

Action Required

Change applications. Do not attempt to run TRUNCATE TABLE on an external table.

Alternatively, change applications to alter a table property to set `external.table.purge` to `true` to allow truncation of an external table:

```
ALTER TABLE mytable SET TBLPROPERTIES ('external.table.purge'='true');
```

Hive unsupported interfaces and features

You need to know the interfaces available in HDP or CDH platforms that are not supported.

Unsupported Interfaces

The following interfaces are not supported in CDP Private Cloud Base:

- Druid
- Hcat CLI (however HCatalog is supported)
- Hive CLI (replaced by Beeline)
- Hive View UI feature in Ambari
- Apache Hive Standalone driver
- Renaming Hive databases
- Multiple insert overwrite queries that read data from a source table.
- LLAP
- MapReduce execution engine (replaced by Tez)
- Pig
- S3 for storing tables (available in CDP Public Cloud only)
- Spark execution engine (replaced by Tez)
- Spark thrift server

Spark and Hive tables interoperate using the Hive Warehouse Connector.

- SQL Standard Authorization
- Storage Based Authorization
- Tez View UI feature in Ambari
- WebHCat

You can use Hue in lieu of Hive View.

Storage Based Authorization

Storage Based Authorization (SBA) is no longer supported in CDP. Ranger integration with Hive metastore provides consistency in Ranger authorization enabled in HiveServer (HS2). SBA did not provide authorization support for metadata that does not have a file/directory associated with it. Ranger-based authorization has no such limitation.

Partially unsupported interfaces

Apache Hadoop Distributed Copy (DistCP) is not supported for copying Hive ACID tables.

Unsupported Features

CDP does not support the following features that were available in HDP and CDH platforms:

- Replicate Hive ACID tables between CDP Private Cloud Base clusters using REPL commands

You cannot use REPL commands (REPL DUMP and REPL LOAD) to replicate Hive ACID tables between CDP Private Cloud Base clusters that are on a version lower than CDP Private Cloud Base 7.1.8. To use REPL commands, ensure that the source cluster is on CDP Private Cloud Base 7.1.8 or a higher version.

- **CREATE TABLE** that specifies a managed table location

Do not use the **LOCATION** clause to create a managed table. Hive assigns a default location in the warehouse to managed tables. That default location is configured in Hive using the `hive.metastore.warehouse.dir` configuration property, but can be overridden for the database by setting the **CREATE DATABASE MANAGEDLOCATION** parameter.

- **CREATE INDEX** and related index commands were removed in Hive 3, and consequently are not supported in CDP.

In CDP, you use the Hive 3 default ORC columnar file formats to achieve the performance benefits of indexing. Materialized Views with automatic query rewriting also improves performance. Indexes migrated to CDP are preserved but render any Hive tables with an undroppable index. To drop the index, google the Known Issue for CDPD-23041.

- Hive metastore (HMS) high availability (HA) load balancing in CDH

You need to set up HMS HA as described in the documentation.

- Local or Embedded Hive metastore server

CDP does not support the use of a local or embedded Hive metastore setup.

Unsupported Connector Use

CDP does not support the Sqoop exports using the Hadoop jar command (the Java API) that Teradata documents. For more information, see [Migrating data using Sqoop](#).

Supplemental Upgrade Topics

Additional topics to help with special situations.

Configuring a Local Package Repository

You can create a package repository for Cloudera Manager either by hosting an internal web repository or by manually copying the repository files to the Cloudera Manager Server host for distribution to Cloudera Manager Agent hosts.



Note: This internal web repository can also be used when your Cloudera Manager server or cluster does not have access to internet. You must download the installable separately from archive.cloudera.com and place the installable in the internal repository.

Loading Filters ... 7.11.3 7.7.3 7.7.1 7.6.7 7.6.1 7.5.1 7.4.4 7.3.1 7.2.4 7.1.4 7.1.3 7.1.2 7.1.1 7.0.3 5.16 5.15 5.14 5.13 5.16 5.15 5.14 5.13 7.1.9 7.1.8 7.1.7.3000 7.1.7.2000 7.1.7.1000 7.1.7 7.1.6 7.1.9.1000 7.1.9 7.1.8 7.1.7.3000 7.1.7.2000 7.1.7.1000 7.1.7



Important: Select a supported operating system for the versions of Cloudera Manager or CDH that you are downloading. See [CDH and Cloudera Manager Supported Operating Systems](#).



Note: [Cloudera Manager 7.7.3](#) should only be used when you need to use Python 3.8 for the Cloudera Manager agents. You must install Python 3.8 on all hosts before installing or upgrading to Cloudera Manager 7.7.3. Cloudera Manager 7.7.3-CHF4 supports only RHEL 8.4, RHEL 8.6, RHEL 7.9, and SLES 15 SP4. See the [CDP Private Cloud Base Installation Guide](#) for more information.



Warning: Upgrades from Cloudera Manager 5.12 and lower to Cloudera Manager 7.1.1 or higher are not supported



Important: Upgrading Cloudera Manager to version 7.7.1 or higher from clusters where CDH 5.x is deployed is not supported. To upgrade such clusters:

1. Upgrade Cloudera Manager to version 6.3.4.
2. Upgrade CDH to version 6.3.4
3. Upgrade Cloudera Manager to version 7.6.5 or higher



Warning: For upgrades from CDH clusters with Cloudera Navigator to Cloudera Runtime 7.1.1 (or higher) clusters where Navigator is to be migrated to Apache Atlas, the cluster must have Kerberos enabled before upgrading.



Warning: Before upgrading CDH 5 clusters with Sentry to Cloudera Runtime 7.1.x clusters where Sentry privileges are to be transitioned to Apache Ranger:

- The cluster must have Kerberos enabled.
- Verify that HDFS gateway roles exist on the hosts that runs the Sentry service.



Important: If HDFS ACL sync is enabled (`hdfs_sentry_sync_enable=true`) on the CDH cluster, then you must install Ranger RMS to support the same functionality. For steps to install Ranger RMS, see [Installing Ranger RMS](#).



Note: If the cluster you are upgrading will include Atlas, Ranger, or both, the upgrade wizard deploys one infrastructure Solr service to provide a search capability of the audit logs through the Ranger Admin UI and/or to store and serve Atlas metadata. Cloudera recommends that you do not use this service for customer workloads to avoid interference with audit and timeline performance.

Creating a Permanent Internal Repository

The following sections describe how to create a permanent internal repository using Apache HTTP Server:

Setting Up a Web server

To host an internal repository, you must install or use an existing Web server on an internal host that is reachable by the Cloudera Manager host, and then download the repository files to the Web server host. The examples in this section use Apache HTTP Server as the Web server. If you already have a Web server in your organization, you can skip to [Downloading and publishing the package repository for Cloudera Manager](#) on page 58.

1. Install Apache HTTP Server:

RHEL / CentOS

```
sudo yum install httpd
```

SLES

```
sudo zypper install httpd
```

Ubuntu

```
sudo apt-get install httpd
```

2. Start Apache HTTP Server:

RHEL 7, 8

```
sudo systemctl start httpd
```

SLES 12, Ubuntu 16 or later

```
sudo systemctl start apache2
```

Downloading and publishing the package repository for Cloudera Manager

1. Download the package repository for the product you want to install:

Cloudera Manager 7

Do the following steps to download the files for a Cloudera Manager release:

- a. Run the following command to create a local repository directory to hold the Cloudera package repository:

```
sudo mkdir -p /var/www/html/cloudera-repos/cm7
```

- b. Run the following command to download the repository tarball for your operating system:

```
wget https://[username]:[password]@archive.cloudera.com/p/cm7/7.0.3/repo-as-tarball/cm7.0.3-redhat7.tar.gz
```

- c. Run the following command to unpack the tarball into the local repository directory:

```
tar xvfz cm7.0.3-redhat7.tar.gz -C /var/www/html/cloudera-repos/cm7 --strip-components=1
```

- d. Run the following command to modify the file permission that allows you to download the files under the local repository directory:

```
sudo chmod -R ugo+rX /var/www/html/cloudera-repos/cm7
```

2. Visit the Repository URL `http://<web_server>/cloudera-repos/` in your browser and verify the files you downloaded are present.



Important: If you do not see the list of downloaded files in your web browser, then you might have been configured not to display indexes. Verify your web browser settings.



Important:

If you downloaded the rpm files individually instead of using `repo-as-tarball`, ensure that the `allkeys*.asc` files (`allkeys.asc` and the `allkeysha256.asc` for Cloudera Manager 7.11.2 and later versions) are present at the top level of the package repository. The `allkeys*.asc` files are included in the `repo-as-tarball` file. Ensure to include `allkeys*.asc` files, if you are manually copying the package files between hosts.

The `allkeys*.asc` files are used to validate the signatures of the package files during host installation. If `allkeys*.asc` files are not available in the repository, then you cannot add a host in the Cloudera Manager.

Creating a Temporary Internal Repository

You can quickly create a temporary remote repository to deploy packages on a one-time basis. Cloudera recommends using the same host that runs Cloudera Manager, or a gateway host. This example uses [Python SimpleHTTPServer](#) as the Web server to host the `/var/www/html` directory, but you can use a different directory.

1. Download the repository you need following the instructions in [Downloading and publishing the package repository for Cloudera Manager](#) on page 58.
2. Determine a port that your system is not listening on. This example uses port 8900.
3. Start a Python SimpleHTTPServer in the `/var/www/html` directory:

```
cd /var/www/html
python -m SimpleHTTPServer 8900
```

```
Serving HTTP on 0.0.0.0 port 8900 ...
```

4. Visit the Repository URL `http://<web_server>:8900/cloudera-repos/` in your browser and verify the files you downloaded are present.

Configuring Hosts to Use the Internal Repository

After establishing the repository, modify the client configuration to use it:

OS	Procedure
RHEL compatible	<p>Create /etc/yum.repos.d/cloudera-repo.repo files on cluster hosts with the following content, where <code><web_server></code> is the hostname of the Web server:</p> <pre>[cloudera-repo] name=cloudera-repo baseurl=http://<web_server>/cm/5 enabled=1 gpgcheck=0</pre>
SLES	<p>Use the zypper utility to update client system repository information by issuing the following command:</p> <pre>zypper addrepo http://<web_server>/cm <alias></pre>
Ubuntu	<p>Create /etc/apt/sources.list.d/cloudera-repo.list files on all cluster hosts with the following content, where <code><web_server></code> is the hostname of the Web server:</p> <pre>deb http://<web_server>/cm <codename> <components></pre> <p>You can find the <code><codename></code> and <code><components></code> variables in the <code>/conf/distributions</code> file in the repository. After creating the <code>.list</code> file, run the following command:</p> <pre>sudo apt-get update</pre>

Configuring a Local Parcel Repository

You can create a parcel repository for Cloudera Runtime either by hosting an internal Web repository or by manually copying the repository files to the Cloudera Manager Server host for distribution to Cloudera Manager Agent hosts.

Loading Filters ... 7.11.3 7.7.3 7.7.1 7.6.7 7.6.1 7.5.1 7.4.4 7.3.1 7.2.4 7.1.4 7.1.3 7.1.2 7.1.1 7.0.3 5.16 5.15 5.14 5.13 5.16 5.15 5.14 5.13 7.1.9 7.1.8 7.1.7.3000 7.1.7.2000 7.1.7.1000 7.1.7 7.1.6 7.1.9.1000 7.1.9 7.1.8 7.1.7.3000 7.1.7.2000 7.1.7.1000 7.1.7

Using an Internally Hosted Remote Parcel Repository

The following sections describe how to use an internal Web server to host a parcel repository:

Setting Up a Web Server

To host an internal repository, you must install or use an existing Web server on an internal host that is reachable by the Cloudera Manager host, and then download the repository files to the Web server host. The examples on this page use Apache HTTP Server as the Web server. If you already have a Web server in your organization, you can skip to [Downloading and Publishing the Parcel Repository](#) on page 62.

1. Install Apache HTTP Server:

RHEL / CentOS


```
sudo yum install httpd
```

SLES

```
sudo zypper install httpd
```

Ubuntu

```
sudo apt-get install httpd
```

2.  **Warning:** Skipping this step could result in an error message Hash verification failed when trying to download the parcel from a local repository, especially in Cloudera Manager 6 and higher.

Edit the Apache HTTP Server configuration file (/etc/httpd/conf/httpd.conf by default) to add or edit the following line in the <IfModule mime_module> section:

```
AddType application/x-gzip .gz .tgz .parcel
```

If the <IfModule mime_module> section does not exist, you can add it in its entirety as follows:



Note: This example configuration was modified from the default configuration provided after installing Apache HTTP Server on RHEL 7.

```
<IfModule mime_module>
#
# TypesConfig points to the file containing the list of mappings from
# filename extension to MIME-type.
#
TypesConfig /etc/mime.types
#
# AddType allows you to add to or override the MIME configuration
# file specified in TypesConfig for specific file types.
#
#AddType application/x-gzip .tgz
#
# AddEncoding allows you to have certain browsers uncompress
# information on the fly. Note: Not all browsers support this.
#
#AddEncoding x-compress .Z
#AddEncoding x-gzip .gz .tgz
#
# If the AddEncoding directives above are commented-out, then you
# probably should define those extensions to indicate media types:
#
AddType application/x-compress .Z
AddType application/x-gzip .gz .tgz .parcel

#
# AddHandler allows you to map certain file extensions to "handlers":
# actions unrelated to filetype. These can be either built into the se
rver
# or added with the Action directive (see below)
#
# To use CGI scripts outside of ScriptAliased directories:
# (You will also need to add "ExecCGI" to the "Options" directive.)
#
#AddHandler cgi-script .cgi
```

```
# For type maps (negotiated resources):
#AddHandler type-map var

#
# Filters allow you to process content before it is sent to the client
.
#
# To parse .shtml files for server-side includes (SSI):
# (You will also need to add "Includes" to the "Options" directive.)
#
AddType text/html .shtml
AddOutputFilter INCLUDES .shtml
</IfModule>
```

3. Start Apache HTTP Server:

RHEL 7, 8

```
sudo systemctl start httpd
```

SLES 12, Ubuntu 16 or later

```
sudo systemctl start apache2
```

Downloading and Publishing the Parcel Repository

1. Look up the *Cloudera Runtime version* number for your deployment on the [Cloudera Runtime Download Information](#) page. You will need this version number in the next step.
2. Download manifest.json and the parcel files for the product you want to install:

To download the files for the latest Runtime 7 release, run the following commands on the Web server host:

```
sudo mkdir -p /var/www/html/cloudera-repos
sudo wget --recursive --no-parent --no-host-directories https://
[username]:[password]@archive.cloudera.com/p/cdh7/Cloudera Runtime
version/parcels/ -P /var/www/html/cloudera-repos

sudo chmod -R ugo+rX /var/www/html/cloudera-repos/cdh7
```

3. Visit the Repository URL `http://<Web_server>/cloudera-repos/` in your browser and verify the files you downloaded are present. If you do not see anything, your Web server may have been configured to not show indexes.

Configuring Cloudera Manager to Use an Internal Remote Parcel Repository

1. Use one of the following methods to open the parcel settings page:
 - Navigation bar
 - a. Click the parcel icon in the top navigation bar or click Hosts and click the Parcels tab.
 - b. Click the Configuration button.
 - Menu
 - a. Select AdministrationSettings.
 - b. Select CategoryParcels.
2. In the Remote Parcel Repository URLs list, click the addition symbol to open an additional row.
3. Enter the path to the parcel. For example: `http://<web_server>/cloudera-parcels/cdh7/7.2.16.0.0/`
4. Enter a Reason for change, and then click Save Changes to commit the changes.

Using a Local Parcel Repository

To use a local parcel repository, complete the following steps:

1. Open the Cloudera Manager Admin Console and navigate to the Parcels page.
2. Select Configuration and verify that you have a Local Parcel Repository path set. By default, the directory is `/opt/cloudera/parcel-repo`.
3. Remove any Remote Parcel Repository URLs you are not using, including ones that point to Cloudera archives.
4. Add the parcel you want to use to the local parcel repository directory that you specified. For instructions on downloading parcels, see [Downloading and Publishing the Parcel Repository](#) on page 62 above.
5. In the command line, navigate to the local parcel repository directory.
6. Create a SHA1 hash for the parcel you added and save it to a file named `parcel_name.parcel.sha`.

For example, the following command generates a SHA1 hash for the parcel `CDH-6.1.0-1.cdh6.1.0.p0.770702-el7.parcel`:

```
sha1sum CDH-6.1.0-1.cdh6.1.0.p0.770702-el7.parcel | awk '{ print $1 }'
> CDH-6.1.0-1.cdh6.1.0.p0.770702-el7.parcel.sha
```

7. Change the ownership of the parcel and hash files to `cloudera-scm`:

```
sudo chown -R cloudera-scm:cloudera-scm /opt/cloudera/parcel-repo/*
```

8. In the Cloudera Manager Admin Console, navigate to the Parcels page.
9. Click Check for New Parcels and verify that the new parcel appears.
10. Download, distribute, and activate the parcel.

Changes to CDH Hive Tables

As a Data Scientist, Architect, Analyst, or other Hive user you need to locate and use your Apache Hive 3 tables after an upgrade. You also need to understand the changes that occur during the upgrade process. The location of existing tables after a CDH to CDP upgrade does not change. Upgrading CDH to CDP Private Cloud Base converts Hive managed tables to external tables in Hive 3.

About this task

When the upgrade process converts a managed table to external, it sets the table property `external.table.purge` to true. The table is equivalent to a managed table having `purge` set to true in your old CDH cluster.

Managed tables on the HDFS in `/user/hive/warehouse` before the upgrade remain there after the conversion to external. Tables that were external before the upgrade are not relocated. You need to set HDFS policies to access external tables in Ranger, or set up HDFS ACLs.


The upgrade process sets the `hive.metastore.warehouse.dir` property to `/warehouse/tablespace/managed/hive`, designating it the Hive warehouse location for managed tables. New managed tables that you create in CDP are stored in the Hive warehouse. New external tables are stored in the Hive external warehouse `/warehouse/tablespace/external/hive`.

To change the location of the Hive warehouses, you navigate to one of the following menu items in the first step below.

- Hive Action Menu Create Hive Warehouse Directory
- Hive Action Menu Create Hive Warehouse External Directory

Procedure

1. Set up directories for the Hive warehouse directory and Hive warehouse external directory from Cloudera Manager Actions.



2. In Cloudera Manager, click Clusters Hive (the Hive Metastore service) Configuration , and change the hive.metastore.warehouse.dir property value to the path you specified for the new Hive warehouse directory.
3. Change the hive.metastore.warehouse.external.dir property value to the path you specified for the Hive warehouse external directory.
4. Configure Ranger policies or set up ACL permissions to access the directories.

Changes to HDP Hive tables

As a Data Scientist, Architect, Analyst, or other Hive user you need to locate and use your Apache Hive 3 tables after an upgrade. You also need to understand the changes that occur during the upgrade process.

Managed, ACID tables that are not owned by the hive user remain managed tables after the upgrade, but hive becomes the owner.

After the upgrade, the format of a Hive table is the same as before the upgrade. For example, native or non-native tables remain native or non-native, respectively.

After the upgrade, the location of managed tables or partitions do not change under any one of the following conditions:

- The old table or partition directory was not in its default location /apps/hive/warehouse before the upgrade.
- The old table or partition is in a different file system than the new warehouse directory.
- The old table or partition directory is in a different encryption zone than the new warehouse directory.

Otherwise, the upgrade process from HDP to CDP moves managed files to the Hive warehouse /warehouse/tablespace/managed/hive. The upgrade process carries the external files over to CDP with no change in location. By default, Hive places any new external tables you create in /warehouse/tablespace/external/hive. The upgrade process sets the hive.metastore.warehouse.dir property to this location, designating it the Hive warehouse location.

Changes to table references using dot notation

Upgrading to CDP includes the Hive-16907 bug fix, which rejects `db.table` in SQL queries. The dot (.) is not allowed in table names. To reference the database and table in a table name, both must be enclosed in backticks as follows: `db`.`table`.

Changes to ACID properties

Hive 3.x in CDP Private Cloud Base supports transactional and non-transactional tables. Transactional tables have atomic, consistent, isolation, and durable (ACID) properties. In Hive 2.x, the initial version of ACID transaction processing was ACID v1. In Hive 3.x, the mature version of ACID is ACID v2, which is the default table type in CDP Private Cloud Base.

Native and non-native storage formats

Storage formats are a factor in upgrade changes to table types. Hive 2.x and 3.x support the following native and non-native storage formats:

- Native: Tables with built-in support in Hive, such as those in the following file formats:
 - Text
 - Sequence File
 - RC File
 - AVRO File
 - ORC File
 - Parquet File
- Non-native: Tables that use a storage handler, such as the `DruidStorageHandler` or `HBaseStorageHandler`

CDP upgrade changes to HDP table types

The following table compares Hive table types and ACID operations before an upgrade from HDP 2.x and after an upgrade to CDP. The ownership of the Hive table file is a factor in determining table types and ACID operations after the upgrade.

Table 50: HDP 2.x and CDP Table Type Comparison

HDP 2.x				CDP	
Table Type	ACID v1	Format	Owner (user) of Hive Table File	Table Type	ACID v2
External	No	Native or non-native	hive or non-hive	External	No
Managed	Yes	ORC	hive or non-hive	Managed, updatable	Yes
Managed	No	ORC	hive	Managed, updatable	Yes
			non-hive	External, with data delete	No
Managed	No	Native (but non-ORC)	hive	Managed, insert only	Yes
			non-hive	External, with data delete	No
Managed	No	Non-native	hive or non-hive	External, with data delete	No

Transitioning from the Cloudera Manager Embedded PostgreSQL Database Server to an External PostgreSQL Database

Cloudera Manager provides an embedded PostgreSQL database server for demonstration and proof of concept deployments when creating a cluster. To remind users that this embedded database is not suitable for production, Cloudera Manager displays the banner text: "You are running Cloudera Manager in non-production mode, which uses an embedded PostgreSQL database. Switch to using a supported external database before moving into production."

If, however, you have already used the embedded database, and you are unable to redeploy a fresh cluster, then you must migrate to an external PostgreSQL database.



Note: This procedure does not describe how to migrate to a database server other than PostgreSQL. Moving databases from one database server to a different type of database server is a complex process that requires modification of the schema and matching the data in the database tables to the new schema. It is strongly recommended that you engage with Cloudera Professional Services if you wish to perform a transition to an external database server other than PostgreSQL.

Prerequisites

Before migrating the Cloudera Manager embedded PostgreSQL database to an external PostgreSQL database, ensure that your setup meets the following conditions:

- The external PostgreSQL database server is running.
- The database server is configured to accept remote connections.
- The database server is configured to accept user logins using md5.
- No one has manually created any databases in the external database server for roles that will be migrated.



Note: To view a list of databases in the external database server (requires default superuser permission):

```
sudo -u postgres psql -l
```

- All health issues with your cluster have been resolved.

For details about configuring the database server, see [Configuring and Starting the PostgreSQL Server](#).



Important: Only perform the steps in [Configuring and Starting the PostgreSQL Server](#). Do not proceed with the creation of databases as described in the subsequent section.

For large clusters, Cloudera recommends running your database server on a dedicated host. Engage Cloudera Professional Services or a certified database administrator to correctly tune your external database server.

Identify Roles that Use the Embedded Database Server

Before you can migrate to another database server, you must first identify the databases using the embedded database server. When the Cloudera Manager Embedded Database server is initialized, it creates the Cloudera Manager database and databases for roles in the Management Services. The Installation Wizard (which runs automatically the first time you log in to Cloudera Manager) or Add Service action for a cluster creates additional databases for roles when run. It is in this context that you identify which roles are used in the embedded database server.

To identify which roles are using the Cloudera Manager embedded database server:

1. Obtain and save the cloudera-scm superuser password from the embedded database server. You will need this password in subsequent steps:

```
head -1 /var/lib/cloudera-scm-server-db/data/generated_password.txt
```

2. Make a list of all services that are using the embedded database server. Then, after determining which services are not using the embedded database server, remove those services from the list. The scm database must remain in your list. Use the following table as a guide:

Table 51: Cloudera Manager Embedded Database Server Databases

Service	Role	Default Database Name	Default Username
Cloudera Manager Server		scm	scm
Cloudera Management Service	Activity Monitor	amon	amon
Hive	Hive Metastore Server	hive	hive
Hue	Hue Server	hue	7uu7uu7uhue

Service	Role	Default Database Name	Default Username
Cloudera Management Service	Navigator Audit Server	nav	nav
Cloudera Management Service	Navigator Metadata Server	navms	navms
Oozie	Oozie Server	oozie_oozie_server	oozie_oozie_server
Cloudera Management Service	Reports Manager	rman	rman
Sentry	Sentry Server	sentry	sentry

3. Verify which roles are using the embedded database. Roles using the embedded database server always use port 7432 (the default port for the embedded database) on the Cloudera Manager Server host.

For Cloudera Management Services:

- a. Select Cloudera Management Service > Configuration, and type "7432" in the Search field.
- b. Confirm that the hostname for the services being used is the same hostname used by the Cloudera Manager Server.



Note:

If any of the following fields contain the value "7432", then the service is using the embedded database:

- Activity Monitor
- Navigator Audit Server
- Navigator Metadata Server
- Reports Manager

For the Oozie Service:

- a. Select Oozie service > Configuration, and type "7432" in the Search field.
- b. Confirm that the hostname is the Cloudera Manager Server.

For Hive, Hue, and Sentry Services:

- a. Select the specific service > Configuration, and type "database host" in the Search field.
 - b. Confirm that the hostname is the Cloudera Manager Server.
 - c. In the Search field, type "database port" and confirm that the port is 7432.
 - d. Repeat these steps for each of the services (Hive, Hue and Sentry).
4. Verify the database names in the embedded database server match the database names on your list (Step 2). Databases that exist on the database server and not used by their roles do not need to be migrated. This step is to confirm that your list is correct.



Note: Do not add the postgres, template0, or template1 databases to your list. These are used only by the PostgreSQL server.

```
psql -h localhost -p 7432 -U cloudera-scm -l
```

```
Password for user cloudera-scm: <password>
```

Name		Owner	List of databases		
	Access		Encoding	Collate	Ctype
amon		amon	UTF8	en_US.UTF8	en_US.U
TF8					
hive		hive	UTF8	en_US.UTF8	en_US.UT
F8					
hue		hue	UTF8	en_US.UTF8	en_US
.UTF8					

```

nav | nav | UTF8 | en_US.UTF8 | en_US.
UTF8 |
navms | navms | UTF8 | en_US.UTF8 | en_US.U
TF8 |
oozie_oozie_server | oozie_oozie_server | UTF8 | en_US.UTF8 | en_US.UT
F8 |
postgres | cloudera-scm | UTF8 | en_US.UTF8 | en_US
.UTF8 |
rman | rman | UTF8 | en_US.UTF8 | en_US.
UTF8 |
scm | scm | UTF8 | en_US.UTF8 | en_US.U
TF8 |
sentry | sentry | UTF8 | en_US.UTF8 | en_US.UT
F8 |
template0 | cloudera-scm | UTF8 | en_US.UTF8 | en_US
.UTF8 | =c/"cloudera-scm"
template1 | cloudera-scm | UTF8 | en_US.UTF8 | en_US.UT
F8 | =c/"cloudera-scm"
(12 rows)

```

You should now have a list of all roles and database names that use the embedded database server, and are ready to proceed with the transition of databases from the embedded database server to the external PostgreSQL database server.

Migrate Databases from the Embedded Database Server to the External PostgreSQL Database Server

While performing this procedure, ensure that the Cloudera Manager Agents remain running on all hosts. Unless otherwise specified, when prompted for a password use the cloudera-scm password.



Note: After completing this transition, you cannot delete the cloudera-scm postgres superuser unless you remove the access privileges for the migrated databases. Minimally, you should change the cloudera-scm postgres superuser password.

1. In Cloudera Manager, stop the cluster services identified as using the embedded database server (see [Identify Roles that Use the Embedded Database Server](#) on page 66). Be sure to stop the Cloudera Management Service as well. Also be sure to stop any services with dependencies on these services. The remaining services will continue to run without downtime.



Note: If you do not stop the services from within Cloudera Manager before stopping Cloudera Manager Server from the command line, they will continue to run and maintain a network connection to the embedded database server. If this occurs, then the embedded database server will ignore any command line stop commands (Step 2) and require that you manually kill the process, which in turn causes the services to crash instead of stopping cleanly.

2. Navigate to Hosts > All Hosts, and make note of the number of roles assigned to hosts. Also take note whether or not they are in a commissioned state. You will need this information later to validate that your scm database was migrated correctly.
3. Stop the Cloudera Manager Server. To stop the server:

```
sudo service cloudera-scm-server stop
```

4. Obtain and save the embedded database superuser password (you will need this password in subsequent steps) from the generated_password.txt file:

```
head -1 /var/lib/cloudera-scm-server-db/data/generated_password.txt
```

5. Export the PostgreSQL user roles from the embedded database server to ensure the correct users, permissions, and passwords are preserved for database access. Passwords are exported as an md5sum and are not visible in plain text. To export the database user roles (you will need the cloudera-scm user password):

```
pg_dumpall -h localhost -p 7432 -U cloudera-scm -v --roles-only -f "/var/
tmp/cloudera_user_roles.sql"
```

6. Edit /var/tmp/cloudera_user_roles.sql to remove any CREATE ROLE and ALTER ROLE commands for databases not in your list. Leave the entries for cloudera-scm untouched, because this user role is used during the database import.
7. Export the data from each of the databases on your list you created in [Identify Roles that Use the Embedded Database Server](#) on page 66:

```
pg_dump -F c -h localhost -p 7432 -U cloudera-scm [database_name] > /var/
tmp/[database_name]_db_backup-$(date +%m-%d-%Y).dump
```

Following is a sample data export command for the scm database:

```
pg_dump -F c -h localhost -p 7432 -U cloudera-scm scm > /var/tmp/scm_db_
backup-$(date +%m-%d-%Y).dump
```

Password:

8. Stop and disable the embedded database server:

```
service cloudera-scm-server-db stop
chkconfig cloudera-scm-server-db off
```

Confirm that the embedded database server is stopped:

```
netstat -at | grep 7432
```

9. Back up the Cloudera Manager Server database configuration file:

```
cp /etc/cloudera-scm-server/db.properties /etc/cloudera-scm-server/db.pr
operties.embedded
```

10. Copy the file /var/tmp/cloudera_user_roles.sql and the database dump files from the embedded database server host to /var/tmp on the external database server host:

```
cd /var/tmp
scp cloudera_user_roles.sql *.dump <user>@<postgres-server>:/var/tmp
```

11. Import the PostgreSQL user roles into the external database server.

The external PostgreSQL database server superuser password is required to import the user roles. If the superuser role has been changed, you will be prompted for the username and password.



Note: Only run the command that applies to your context; do not execute both commands.

- To import users when using the default PostgreSQL superuser role:

```
sudo -u postgres psql -f /var/tmp/cloudera_user_roles.sql
```

- To import users when the superuser role has been changed:

```
psql -h <database-hostname> -p <database-port> -U <superuser> -f /var/tmp/cloudera_user_roles.sql
```

For example:

```
psql -h pg-server.example.com -p 5432 -U postgres -f /var/tmp/cloudera_user_roles.sql
```

```
Password for user postgres
```

12. Import the Cloudera Manager database on the external server. First copy the database dump files from the Cloudera Manager Server host to your external PostgreSQL database server, and then import the database data:

Note: To successfully run the `pg_restore` command, there must be an existing database on the database server to complete the connection; the existing database will not be modified. If the `-d <existing-database>` option is not included, then the `pg_restore` command will fail.

```
pg_restore -C -h <database-hostname> -p <database-port> -d <existing-database> -U cloudera-scm -v <data-file>
```

Repeat this import for each database.

The following example is for the scm database:

```
pg_restore -C -h pg-server.example.com -p 5432 -d postgres -U cloudera-scm -v /var/tmp/scm_server_db_backup-20180312.dump
```

```
pg_restore: connecting to database for restore
Password:
```

13. Update the Cloudera Manager Server database configuration file to use the external database server. Edit the `/etc/cloudera-scm-server/db.properties` file as follows:

- Update the `com.cloudera.cmf.db.host` value with the hostname and port number of the external database server.
- Change the `com.cloudera.cmf.db.setupType` value from "EMBEDDED" to "EXTERNAL".

14. Start the Cloudera Manager Server and confirm it is working:

```
service cloudera-scm-server start
```

Note that if you start the Cloudera Manager GUI at this point, it may take up to five minutes after executing the start command before it becomes available.

In Cloudera Manager Server, navigate to Hosts > All Hosts and confirm the number of roles assigned to hosts (this number should match what you found in Step 2); also confirm that they are in a commissioned state that matches what you observed in Step 2.

15. Update the role configurations to use the external database hostname and port number. Only perform this task for services where the database has been migrated.

For Cloudera Management Services:

- a. Select Cloudera Management Service > Configuration, and type "7432" in the Search field.
- b. Change any database hostname properties from the embedded database to the external database hostname and port number.
- c. Click Save Changes.

For the Oozie Service:

- a. Select Oozie service > Configuration, and type "7432" in the Search field.
- b. Change any database hostname properties from the embedded database to the external database hostname and port number.
- c. Click Save Changes.

For Hive, Hue, and Sentry Services:

- a. Select the specific service > Configuration, and type "database host" in the Search field.
- b. Change the hostname from the embedded database name to the external database hostname.
- c. Click Save Changes.

16. Start the Cloudera Management Service and confirm that all management services are up and no health tests are failing.

17. Start all Services via the Cloudera Manager web UI. This should start all services that were stopped for the database transition. Confirm that all services are up and no health tests are failing.

18. On the embedded database server host, remove the embedded PostgreSQL database server:

- a. Make a backup of the /var/lib/cloudera-scm-server-db/data directory:

```
tar czvf /var/tmp/embedded_db_data_backup-$(date +%m-%d-%Y).tgz /var/lib/cloudera-scm-server-db/data
```

- b. Remove the embedded database package:

For RHEL/SLES:

```
rpm --erase cloudera-manager-server-db-2
```

For Ubuntu:

```
apt-get remove cloudera-manager-server-db-2
```

- c. Delete the /var/lib/cloudera-scm-server-db/data directory.