

CDP Private Cloud Base 7

# In-Place Upgrade HDP 2 to CDP Private Cloud Base

Date published: 2019-11-22

Date modified: 2024-12-10

The Cloudera logo is displayed in a bold, orange, sans-serif font. The word "CLOUDERA" is written in all caps, with the letter 'E' in "CLouDERA" featuring a stylized horizontal line through its center.

<https://docs.cloudera.com/>

# Legal Notice

© Cloudera Inc. 2025. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

# Contents

|  |           |
|--|-----------|
| <b>HDP to CDP Upgrade Overview.....</b>                      | <b>8</b>  |
| In-place upgrade overview.....                               | 8         |
| CDP Upgrade Readiness.....                                   | 13        |
| How much time should I plan for to complete my upgrade?..... | 14        |
| <br>   |           |
| <b>Cluster environment readiness.....</b>                    | <b>15</b> |
| Disk space and mountpoint considerations.....                | 15        |
| Downloading and Publishing the Package Repository.....       | 17        |
| Downloading and Publishing the Parcel Repository.....        | 20        |
| Hadoop Users (user:group) and Kerberos Principals.....       | 22        |
| <br>   |           |
| <b>Sample data ingestion.....</b>                            | <b>38</b> |
| <br>   |           |
| <b>Expediting the Hive upgrade.....</b>                      | <b>39</b> |
| Overview of the expedited Hive upgrade.....                  | 39        |
| Handle Missing Table or Partition Locations.....             | 40        |
| Managed Table Location Mapping.....                          | 40        |
| Check SERDE Definitions and Availability.....                | 40        |
| Make Tables SparkSQL Compatible.....                         | 40        |
| Understanding the Hive upgrade.....                          | 41        |
| <br>   |           |
| <b>Modifying the HSMM to prevent migration.....</b>          | <b>41</b> |
| <br>   |           |
| <b>Ambari and HDP Upgrade Checklist.....</b>                 | <b>42</b> |
| Ambari upgrade checklist.....                                | 43        |
| Download cluster blueprints without hosts.....               | 43        |
| HDP upgrade checklist.....                                   | 44        |
| Checklist for large clusters.....                            | 44        |
| Kerberos cluster.....  | 45        |
| Before upgrading any cluster.....                            | 45        |
| Managing MPacks.....   | 46        |
| Changes to Ambari and HDP services.....                      | 47        |
| HDP Core component version changes.....                      | 48        |
| Upgrading the cluster's underlying OS.....                   | 48        |
| In-Place and Restore.....                                    | 49        |
| Move and Decommission.....                                   | 49        |
| <br>   |           |
| <b>Upgrading Ambari.....</b>                                 | <b>49</b> |
| Before you upgrade Ambari.....                               | 50        |
| Setting up a local repository.....                           | 51        |
| Updating Ambari repo files.....                              | 51        |
| Updating HDP repo files.....                                 | 51        |
| Case study for setting up an HDP-GPL local repository.....   | 51        |

|   |    |
|---|----|
| Setting up local repository with temporary internet access.....             | 52 |
| Case study for setting up local repository.....                             | 53 |
| Update version repository base urls.....                                    | 53 |
| Preparing Ambari Repository Configuration File to use Local Repository..... | 54 |
| Backup Ambari.....  | 54 |
| Ambari Behavioral changes.....  | 55 |
| Ambari Properties backup.....   | 56 |
| Review Ambari UI and the Quick Links.....                                   | 57 |
| Upgrade to Ambari 7.1.x.0.....  | 57 |
| Upgrade to Ambari 7.1.7 SP2 (7.1.7.2000-3).....                             | 57 |
| Upgrade to Ambari 7.1.7.1.....  | 62 |
| Upgrade to Ambari 7.1.6.0.....  | 67 |
| Upgrade to Ambari 7.1.4.0.....  | 72 |
| Upgrade to Ambari 7.1.3.0.....  | 77 |
| Upgrade to Ambari 7.1.2.0.....  | 81 |
| Upgrade to Ambari 7.1.0.0.....  | 85 |
| Download cluster blueprints.....  | 89 |
| Mandatory Post-Upgrade Tasks.....   | 90 |
| Upgrading Ambari Metrics System and SmartSense.....                         | 90 |

## **Upgrading HDP to Cloudera Runtime 7.1.x.....93**

|   |     |
|---|-----|
| HDP Prerequisites.....  | 93  |
| Upgrade process.....  | 93  |
| Kerberos cluster.....   | 94  |
| Before upgrading any cluster.....                               | 94  |
| Backup HDP Cluster.....   | 95  |
| Backup and Restore Databases.....                               | 95  |
| Backup Ranger.....  | 96  |
| Backup Atlas.....   | 97  |
| Backup Ambari Infra Solr.....                                   | 98  |
| Backup Hive.....  | 99  |
| Backup HBase.....   | 100 |
| Backup Kafka.....   | 100 |
| Backup Oozie.....   | 100 |
| Backup Knox.....  | 101 |
| Backup Logsearch.....   | 101 |
| Backup Zeppelin.....  | 101 |
| Backup HDFS.....  | 101 |
| Backup ZooKeeper.....   | 102 |
| Backup Ambari.....  | 103 |
| Backup databases.....   | 103 |
| Before you upgrade.....   | 103 |
| Register software repositories.....                             | 104 |
| HDP Intermediate bits for 7.1.x.0 Repositories.....             | 105 |
| Software download matrix for HDP 2.6.5 to CDP 7.1.x.....        | 133 |
| Install software on the hosts.....                              | 134 |
| Preparing the services for upgrade.....                         | 135 |
| Backing up Ambari infra data.....                               | 135 |
| Overview of the Migration of the Atlas and Infra Solr Data..... | 139 |
| Ranger Service connection with Oracle database.....             | 141 |
| Ranger admin password.....                                      | 141 |
| Preparing HBase for upgrade.....                                | 142 |
| Preparing HDP Search for upgrade.....                           | 146 |
| Preparing Hive for upgrade.....                                 | 156 |
| Preparing Kafka for upgrade.....                                | 162 |

|   |     |
|---|-----|
| Preparing Zeppelin for upgrade.....             | 163 |
| Perform the HDP upgrade.....                    | 163 |
| Perform express upgrade.....                    | 164 |
| Post-HDP-upgrade tasks.....                     | 165 |
| Update Ranger passwords.....                    | 165 |
| Atlas Migration and HBase Hook settings.....    | 166 |
| Ambari Metrics and LogSearch.....               | 166 |
| Ambari infra-migrate and restore.....           | 166 |
| Upload HDFS entity information.....             | 167 |
| Custom Spark SQL Warehouse Directory.....       | 167 |
| Hive post-HDP-upgrade tasks.....                | 167 |
| Verify Zeppelin settings in Ambari.....         | 171 |
| Search post-HDP-upgrade tasks.....              | 172 |
| Backup Infra Solr collections.....              | 172 |
| Troubleshooting the HDP upgrade.....            | 173 |
| Hive Metastore corrupt.....                     | 173 |
| Missing Hive tables.....                        | 174 |
| YARN Registry DNS instance fails to start.....  | 174 |
| Ambari Metrics System (AMS) does not start..... | 174 |
| Ranger MySQL collation.....                     | 175 |
| Rollback Ambari to 2.6.5.....                   | 180 |
| Rollback HDP Services.....                      | 181 |
| Overview.....                                   | 181 |
| ZooKeeper.....                                  | 182 |
| Ambari-Metrics.....                             | 182 |
| Ambari Infra Solr.....                          | 183 |
| Ranger.....                                     | 184 |
| HDFS.....                                       | 187 |
| YARN.....                                       | 189 |
| HBase.....                                      | 189 |
| Kafka.....                                      | 189 |
| Atlas.....                                      | 189 |
| Hive.....                                       | 190 |
| Spark.....                                      | 190 |
| Oozie.....                                      | 190 |
| Knox.....                                       | 190 |
| Zeppelin.....                                   | 190 |
| Log Search.....                                 | 190 |

## **Transitioning to Cloudera Manager..... 191**

|   |     |
|---|-----|
| Pre-transition steps.....   | 192 |
| Databases.....  | 193 |
| Kerberos.....   | 193 |
| Atlas migration.....  | 193 |
| HDFS.....   | 194 |
| Spark.....  | 194 |
| Ranger.....   | 195 |
| Kerberos - Optional task.....   | 195 |
| Solr.....   | 195 |
| Cloudera Manager Installation and Setup.....  | 195 |
| Installing JDBC Driver.....   | 197 |
| Proxy Cloudera Manager through Apache Knox.....   | 199 |
| Transitioning HDP cluster to CDP Private Cloud Base cluster using the AM2CM tool.....             | 200 |
| Transitioning HDP 2.6.5 cluster to CDP Private Cloud Base 7.1.x cluster using the AM2CM tool..... | 200 |

|  |     |
|--|-----|
| Post transition steps.....                                       | 205 |
| Enable Auto Start setting.....                                   | 205 |
| Kerberos Principal for Cloudera Manager Server.....              | 205 |
| ZooKeeper.....   | 205 |
| Delete ZNODES.....   | 205 |
| Ranger.....  | 206 |
| Ranger KMS.....  | 206 |
| Add Ranger policies for components on the CDP Cluster.....       | 207 |
| Ranger Installation in High Availability with Load Balancer..... | 211 |
| Set maximum retention days for Ranger audits.....                | 212 |
| HDFS.....  | 212 |
| Solr.....  | 214 |
| Kafka.....   | 216 |
| YARN.....  | 217 |
| Spark.....   | 222 |
| Tez.....   | 225 |
| Hive.....  | 227 |
| HBase.....   | 228 |
| HBase RegionServer heap size.....                                | 229 |
| Hue.....   | 229 |
| Oozie.....   | 243 |
| Atlas advanced configuration snippet (Safety valve).....         | 246 |
| Migrating Atlas data.....  | 246 |
| Phoenix.....   | 248 |
| Starting all services.....                                       | 249 |
| Hive Policy Additions.....                                       | 249 |
| Knox.....  | 249 |
| Client Configurations.....                                       | 250 |
| Securing ZooKeeper.....  | 251 |
| Zeppelin Shiro configurations.....                               | 251 |
| Migrating Spark workloads to .....                               | 251 |
| Apache Hive Expedited Migration Tasks.....                       | 271 |
| Apache Hive Changes in CDP.....                                  | 273 |
| Apache Hive Post-Upgrade Tasks.....                              | 291 |

## **Configuring External Authentication for Cloudera Manager..... 302**

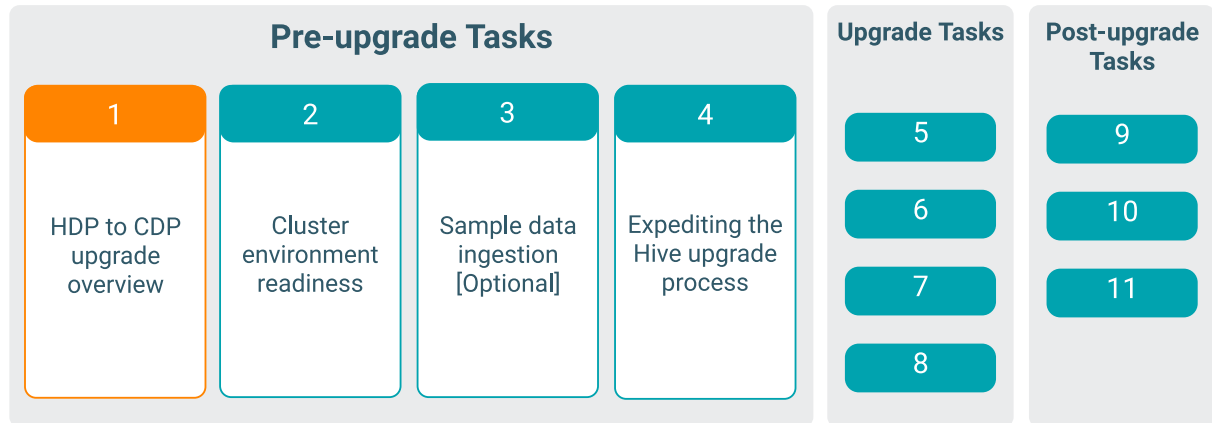
## **Additional Services..... 302**

|   |     |
|---|-----|
| Installing Data Analytics Studio on a Cloudera Private Cloud Base cluster using Ambari.....                                 | 303 |
| Check cluster configuration for Hive and Tez.....   | 303 |
| Add the DAS service.....  | 303 |
| DAS post-installation tasks.....  | 304 |
| Troubleshooting DAS installation.....   | 311 |
| Problem area: Queries page.....   | 311 |
| Problem area: Compose page.....   | 313 |
| Problem area: Reports page.....   | 314 |
| DAS service installation fails with the “python files missing” message.....   | 314 |
| DAS does not log me out as expected, or I stay logged in longer than the time specified in the<br>Ambari configuration..... | 314 |
| Getting a 401 - Unauthorized access error message while accessing DAS.....  | 315 |
| Installing DAS on CDP Private Cloud Base using Cloudera Manager.....  | 316 |
| Adding Hue service with Cloudera Manager.....   | 317 |
| Install and configure MySQL database.....   | 317 |
| Add the Hue service using Cloudera Manager.....   | 317 |

|   |            |
|---|------------|
| Enable Kerberos for authentication.....   | 319        |
| Integrate Hue with Knox.....  | 320        |
| Grant Ranger permissions to new users or groups.....  | 322        |
| Adding Query Processor service to a cluster.....  | 324        |
| <br>  |            |
| <b>Applications Upgrade.....</b>  | <b>325</b> |
| <br>  |            |
| <b>Procedure to Rollback from CDP Private Cloud Base 7.1.7 SP1 to CDP<br/>Private Cloud Base 7.1.7.....</b> | <b>326</b> |

## HDP to CDP Upgrade Overview

The process of upgrading to CDP Private Cloud Base involves understanding the supported in-place upgrade paths and verifying the software and hardware considerations and requirements prior to performing the upgrade steps.



### In-place upgrade overview

To plan your upgrade from Ambari managed HDP 2.6.5 to Cloudera Private Cloud Base, you must be aware of the two stages of in-place upgrade along with the pre-upgrade, upgrade, and post-upgrade tasks.

#### About this task

This upgrade path consists of two stages:

1. Upgrade HDP 2.6.5 to Cloudera Runtime 7.1.x using Ambari.
2. Transition the management platform from Ambari to Cloudera Manager.



**Note:** One stage upgrade from HDP 2.6.5 to CDP Private Cloud Base is now supported using the appropriate CMA version. Cloudera recommends you to upgrade your HDP 2.6.5 cluster to a CDP Private Cloud Base cluster using the One-stage upgrade. For information on the supported upgrade paths and the supported CMA tool, see [In-place upgrade overview](#) documentation.



**Note:** The importance of security in a production environment cannot be understated. TLS and Kerberos form the baseline for secure operations of your CDP Runtime environment. Cloudera supports security services such as Ranger and Atlas only when they are run on clusters where Kerberos is enabled to authenticate users.

To upgrade to Cloudera Runtime 7.1.x, you must first upgrade from Ambari 2.6.2.x (source) to Ambari 7.1.x.x (target) and then upgrade from HDP 2.6.5 to Cloudera Runtime 7.1.x.

Upgrading and transitioning to Cloudera Private Cloud Base overview:



## Upgrading HDP 2.6.5 to CDP

### Pre-upgrade Tasks

- 1 [HDP to CDP Upgrade Overview](#). Review the [Upgrade document topic](#) for the supported upgrade paths.
- 2 [Cluster environment readiness](#). Gather information on your deployment and verify cluster environment readiness.
- 3 [Sample data ingestion \[Optional\]](#). Prepare sample data for ingestion.
- 4 [Expediting the Hive Upgrade process](#). Validate Apache Hive constructs before you upgrade.

### Upgrade Tasks

- 5 [Ambari and HDP upgrade checklist](#). Plan how and when to begin your upgrade.
- 6 [Upgrading Ambari](#). Check readiness and upgrade from Ambari 2.6.2.x to Ambari 7.1.x.
- 7 [Upgrading HDP to Cloudera Runtime 7.1.x](#). Upgrade from HDP 2.6.5 to HDP 7.1.x and complete post-upgrade tasks.
- 8 [Transitioning to Cloudera Manager](#). Transition from HDP to Cloudera Manager.

### Post-upgrade Tasks

- 9 [Configuring external authentication for Cloudera Manager](#). Configure external authentication with Cloudera Manager if you have LDAP configured on your cluster.
- 10 [Additional services](#). Install additional services like Hue on your cluster using Cloudera Manager.
- 11 [Applications upgrade](#). Test and update applications and services on your CDP cluster.

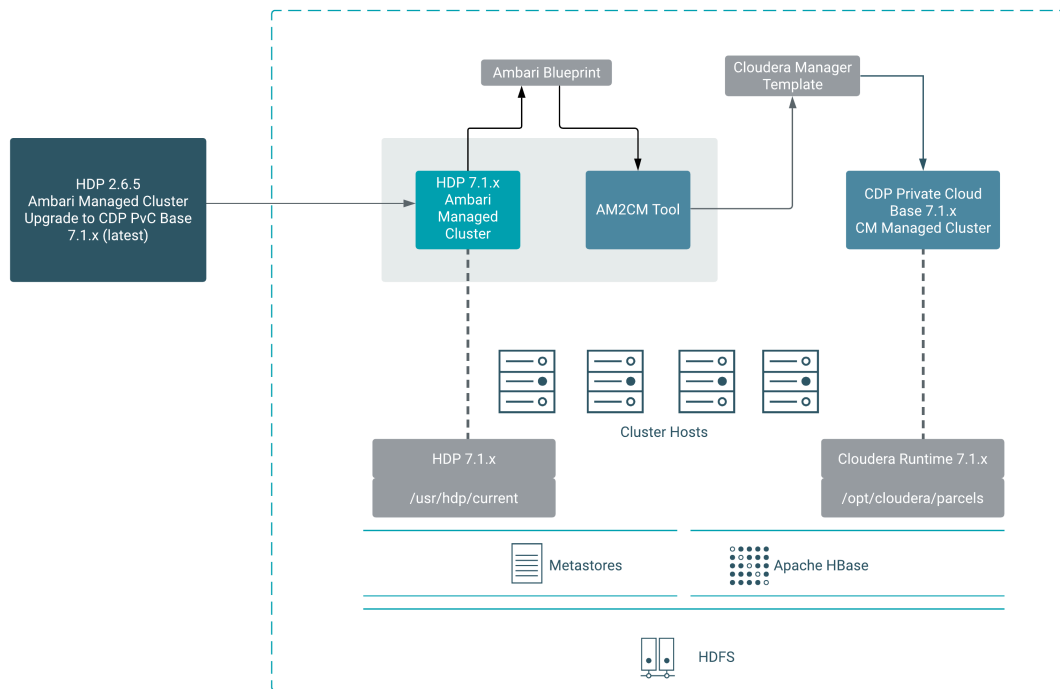


**Note:** If you are upgrading to 7.5.1 or higher in order to install version 1.5.4, you must use version 7.1.6 or higher. For more information, see .



**Note:** If you want to upgrade CDP Private Cloud Data Services, you must upgrade using the following path:

1. Upgrade to Cloudera Manager 7.6.7.
2. Upgrade the Base cluster to 7.1.7 SP2.
3. Upgrade CDP Private Cloud Data Services to 1.5.0.



**Important:**

- While the upgrade process will temporarily install HDP intermediate bits, the HDP intermediate bits is an interim artifact that enables you to complete the upgrade to Cloudera Private Cloud Base and is not a new release of HDP intended for production use. When the upgrade process is complete, your cluster will be using the Cloudera Runtime 7.1.x parcel.
- If you stay on HDP 7.1.x you will not be able to upgrade directly to CDP-7.1.y (where  $y > x$ ). You must always first upgrade from HDP 7.1.x to CDP 7.1.x and then upgrade to CDP-7.1.y.
- The following table shows the supported versions of the AM2CM tool, Cloudera Manager, and Runtime for each upgrade path:

| Upgrade path                             | AM2CM version  | Cloudera Manager       | Runtime   |
|--|--|------------------------|---|
| HDP 2.6.5 to CDP 7.1.7 SP2               | AM2CM 2.8.1.0,<br>AM2CM 2.4.3.0,<br>AM2CM 2.4.2.0,<br>AM2CM 2.4.1.1 and<br>2.4.1.0 | Cloudera Manager 7.6.7 | CDH-7.1.7 SP2   |
| HDP 2.6.5 to CDP 7.1.8                   | AM2CM 2.8.1.0,<br>AM2CM 2.4.3.0,<br>AM2CM 2.4.2.0,<br>AM2CM 2.4.1.1 and<br>2.4.0.0 | Cloudera Manager 7.7.1 | CDH-7.1.8   |
| HDP 2.6.5 to CDP 7.1.7 SP1               | AM2CM 2.8.1.0,<br>AM2CM 2.4.3.0,<br>AM2CM 2.4.2.0,<br>AM2CM 2.4.1.1 and<br>2.4.1.0 | Cloudera Manager 7.6.1 | CDH-7.1.7 SP1   |
| HDP 2.6.5 to CDP 7.1.7 or 7.1.6 or 7.1.4 | AM2CM 2.8.1.0,<br>2.1.0.0, 2.0.4.0, 2.0.3.0,<br>and 2.0.2.0                        | Cloudera Manager 7.4.4 | CDH-7.1.7, CDH-7.1.6,<br>and CDH-7.1.4.6 hotfix<br>versions |
| HDP 2.6.5 to CDP 7.1.7                   | 2.0.0.0  | Cloudera Manager 7.4.4 | CDH-7.1.7   |
| HDP 2.6.5 to CDP 7.1.6                   | 1.1.1.0  | Cloudera Manager 7.3.1 | CDH-7.1.6   |
| HDP 2.6.5 to CDP 7.1.4                   | 1.1.0.0  | CM-7.2.4_patch4455     | CDH-7.1.4.6 hotfix<br>versions                              |

- The following table shows the supported versions of HDP and Ambari on the source cluster, Cloudera Manager and CDP on the target cluster, and the AM2CM tool:

| HDP Version | Operating System               | Ambari Version | AM2CM Version  | Cloudera Manager Version | CDP Version            |
|-------------|--------------------------------|----------------|--|--------------------------|------------------------|
| HDP 2.6.5   | RHEL 7, Sles 12, and Ubuntu 18 | Ambari 2.6.2.x | AM2CM 2.8.1.0,<br>AM2CM 2.4.3.0,<br>AM2CM 2.4.2.0,<br>AM2CM 2.4.1.1 and<br>AM2CM 2.4.1.0 | Cloudera Manager 7.6.7   | CDP PvC Base 7.1.7 SP2 |

|           |                                |                |  |                                   |                        |
|-----------|--------------------------------|----------------|--|-----------------------------------|------------------------|
| HDP 2.6.5 | RHEL 7, Sles 12, and Ubuntu 18 | Ambari 2.6.2.x | AM2CM 2.8.1.0,AM2CM 2.4.3.0, AM2CM 2.4.2.0, AM2CM 2.4.1.1 and AM2CM 2.4.0.0  | Cloudera Manager 7.7.1            | CDP PvC Base 7.1.8     |
| HDP 2.6.5 | RHEL 7, Sles 12, and Ubuntu 18 | Ambari 2.6.2.x | AM2CM 2.8.1.0,AM2CM 2.4.3.0, AM2CM 2.4.2.0, AM2CM 2.4.1.1 and AM2CM 2.4.1.0  | Cloudera Manager 7.6.1            | CDP PvC Base 7.1.7 SP1 |
| HDP 2.6.5 | RHEL 7, Sles 12, and Ubuntu 18 | Ambari 2.6.2.x | AM2CM 2.8.1.0,AM2CM 2.1.0.0, AM2CM 2.0.4.0, AM2CM 2.0.3.0, and AM2CM 2.0.2.0 | Cloudera Manager 7.4.4            | CDP PvC Base 7.1.7     |
| HDP 2.6.5 | RHEL 7, Sles 12, and Ubuntu 18 | Ambari 2.6.2.x | AM2CM 2.8.1.0,AM2CM 2.0.0.0  | Cloudera Manager 7.4.4            | CDP PvC Base 7.1.7     |
| HDP 2.6.5 | RHEL 7                         | Ambari 2.6.2.x | AM2CM 2.8.1.0,AM2CM 1.1.1.0  | Cloudera Manager 7.3.1            | CDP PvC Base 7.1.6     |
| HDP 2.6.5 | RHEL 7                         | Ambari 2.6.2.x | AM2CM 2.8.1.0,AM2CM 1.1.0.0  | Cloudera Manager 7.2.4_patch-4455 | CDP PvC Base 7.1.4     |
| HDP 2.6.5 | RHEL 7                         | Ambari 2.6.2.x | AM2CM 1.0.0.0 (Technical preview)  | Cloudera Manager 7.1.3            | CDP PvC Base 7.1.3     |
| HDP 2.6.5 | RHEL 7                         | Ambari 2.6.2.x | AM2CM 1.0.0.0 (Technical preview)  | Cloudera Manager 7.1.2            | CDP PvC Base 7.1.2     |
| HDP 2.6.5 | RHEL 7                         | Ambari 2.6.2.x | AM2CM 1.0.0.0 (Technical preview)  | Cloudera Manager 7.1.1            | CDP PvC Base 7.1.1     |

The final step of the upgrade process is to transition from the Ambari to Cloudera Manager 7.x.x using the AM2CM tool. It includes the following steps:

- HDP cluster blueprint is exported from Ambari and converted to a Cloudera Manager deployment template using the AM2CM tool.
- Deploy the template to Cloudera Manager 7.x.x and activate the Cloudera Runtime 7.1.x parcels.
- Start using the CDP Private Cloud Base cluster with Cloudera Manager 7.x.x.
- Uninstall Ambari and the HDP stack

**Note:**

Whether you are upgrading to incorporate new features into your environment, or to stay within the lifecycle of products that are serviced and supported, you will eventually want to file an Upgrade Planning Case. When upgrading to CDP Private Cloud Base, Cloudera recommends working with Cloudera Professional Services to simplify your journey and get faster time to value. Cloudera PS offers SmartUpgrade to help you efficiently upgrade or migrate to CDP Private Cloud Base with minimal disruptions to your SLAs. Alternatively, you and your team may have the desire and expertise to perform the upgrade yourself. The choice is yours and Cloudera is here to help.

The [Upgrade Planning Cases Guide](#) will help you file an Upgrade Planning case to:

- Receive suggestions on which version of the CDP Private Cloud runtime to deploy
- Obtain other recommendations to have a better upgrade experience
- Communicate with Cloudera when your expected upgrade will take place
- Obtain a limited scope "Extended Support" for out of support products while planning to Upgrade to CDP-DC. This limited support will be available until June 30, 2021.



**Important:** Hive LLAP and Druid services are not part of Cloudera Private Cloud Base.

If you are running Hive LLAP and Druid workloads, contact Cloudera account team or Cloudera Support for further assistance.



**Note:** Rolling Upgrades are not supported when upgrading to CDP Private Cloud Base.



**Note:** If you want to upgrade from HDP 3.1.5 to CDP Private Cloud Base 7.1.6 or higher, see [HDP 3.1.5 to CDP 7.1.6 or higher](#) documentation.

Troubleshooting: A selection of Cloudera Knowledge Base articles are available that describe common issues encountered by Cloudera customers during upgrades and migrations. See [CDP Upgrade/Migrate Troubleshooting Articles](#). (Cloudera login required.)

## CDP Upgrade Readiness

In preparation for the upgrade, review the cluster environment requirements. You must also be familiar with the new features, behavior changes, and impact of the upgrade on existing configurations before the upgrade.

Your Cloudera Account team can help you assess the impact of the following with respect to your upgrade. Cloudera will work with you to ensure your cluster is ready and meets all the criteria.

- [Cluster environment readiness](#) - latest supported versions
  - Ensure supported Operating System, HDP stack, database and java JDK versions are up-to-date and supported.
  - See the upgrade documentation to ensure that your environments are ready for the upgrade.
- On removed components, services, and functionalities, see [Hive unsupported interfaces and features](#), [HDP Core component version changes](#), and [Changes to Ambari and HDP services](#).
- Changes in behaviour:
  - You need to do a number of migration-related tasks due to semantic changes, and a couple of syntax changes in Hive 3, for example db.table references and DROP CASCADE, might require changes to your applications.
  - 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.

- Downtime:
  - In-place upgrades will require downtime and automation of the tasks will limit the extent of the downtime
  - Migrate to new cluster options may limit downtime.
- HDP Search
  - The version of Solr has moved directly from Solr v. 5.x to v 8.4.1. Solr does not support migrations or upgrades greater than one version and therefore consideration needs to be given on how to solve the upgrade to HDP Search Solr indexes to the latest version.
  - Note that CDP stores its indexes in HDFS unlike HDP Search which typically stores them in files on the local file system.
- Ports - Reset the old ports to standard Cloudera Manager ports. It simplifies the Cloudera Manager managed CDP Private Cloud Base environment for better future support and standards.
- Network multihoming is not supported in CDP. If the HDP cluster is configured for multihomed networks, then you must reconfigure to disable that. For more information, see [Networking and Security Requirements](#).
- Third-party application readiness - Ensure that the third-party software products you are integrating with are certified to work with the HDP intermediate bits and Cloudera Private Cloud Base.
- Bespoke configuration - Consider the impact of upgrading on any bespoke applications and customizations either to or integrating with the HDP platform. For example, devops integration tooling, monitoring, and so on.
- Ambari mpacks are not all compatible with HDP 7.1.x and CDP 7.1.x (CSDs are the Cloudera Manager equivalent of mpacks)
- Not all Management Packs are compatible with the HDP intermediate bits. Cloudera Support can help you assess the feasibility of your upgrade if you have Management Packs other than HDF installed.

## How much time should I plan for to complete my upgrade?

An in-place upgrade can take a variable amount of time to complete. Learn about how to plan for and shorten the amount of time required for your upgrade.

The amount of time required for an in-place upgrade depends on many factors, including:

- The number of hosts in your clusters.
- The mix of services you have deployed in your clusters.
- The amount of data stored in your clusters.

Generally, an upgrade can be completed in 24-48 hours. Upgrades from HDP to may take somewhat longer due to the Ambari to conversion process (AM2CM).

The following table provides some additional information to help you plan for your upgrade.

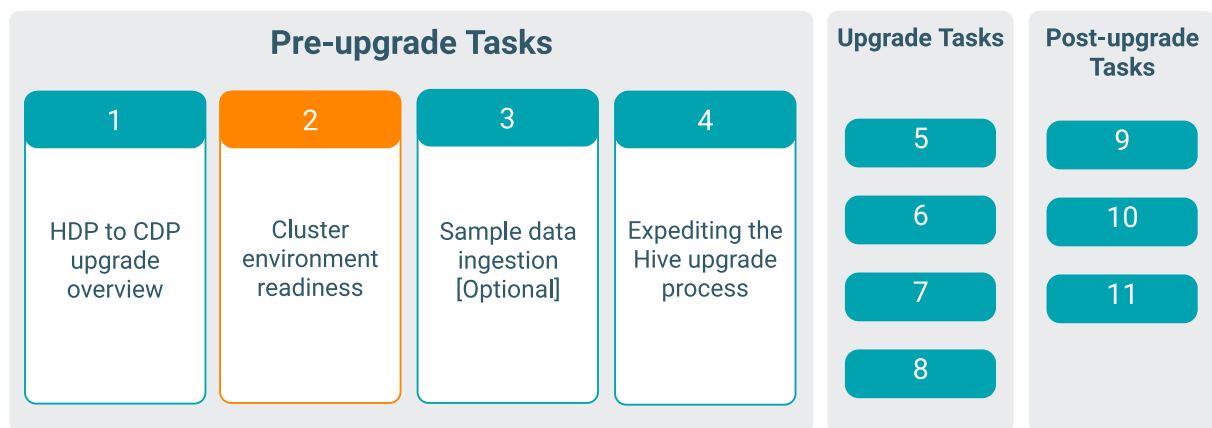
**Table 1: Upgrade Time Planning**

| Component/Process  | Notes  |
|--------------------|--|
| Parcel             | The parcel must be distributed to all hosts before upgrading the hosts. Downloading the parcel directly from archive.cloudera.com over the internet may add additional time. You can download the parcels and serve them from a local web server to decrease this time.<br><br>In addition, after downloading the parcels to a local repository, you can distribute them in advance of launching the upgrade wizard to save additional time. |
|                    | You must upgrade before upgrading your clusters. can continue to manage older versions of and CDH until the upgrade.   |
| Cluster cold start | The cluster will need to be restarted at least once during an in-place upgrade. The amount of time required for a restart depends on how many files and blocks are stored in the cluster and the number of hosts in the cluster.   |
| Hive               | The Hive strict managed migration process can take a significant amount of time. See for more information about mitigating that impact. See <a href="#">Understanding the Hive upgrade (CDH)</a>   |

| Component/Process  | Notes  |
|--------------------|--|
| HBase checks       | While Running HBase checks does not take significant time, remediating any issues can take significant time. To save time during the upgrade, you can plan to do this before running the Upgrade Wizard. |
| Solr export/backup | This process depends on how much data has to be imported after the upgrade.  |

## Cluster environment readiness

You must ensure all the nodes have the supported operating system, Java version, and base Ambari version. Verify the disk space and mount point requirements before you begin the upgrade to the recommended interim HDP bits and CDP Private Cloud Base.




|                                     |  |
|-------------------------------------|--|
| Operating System on all nodes       | <a href="#">Operating System &amp; Upgrading the cluster's underlying OS</a>   |
| Requirements and Supported Versions | <a href="#">CDP Private Cloud Base Requirements and Supported Versions</a>   |
| Java version on all nodes           | <a href="#">Java Versions</a>  |
| Ambari base version on all nodes    | <ul style="list-style-type: none"> <li>For RHEL and CentOS, Ambari version is 2.6.2.x.</li> <li>For Ubuntu, Ambari version is 2.6.2.45.</li> </ul> |
| Repositories                        | <a href="#">Software download matrix</a>   |
| Review disk space                   | <a href="#">Disk space and mountpoint considerations</a>   |

## Disk space and mountpoint considerations

Review the minimum disk space requirements before you upgrade from HDP 2.6.5 to CDP 7.1.x.

| Partition     | Storage                               | Detail  |
|---------------|---------------------------------------|---|
| /usr/hdp      | 10 GB                                 | Minimum space required for each installed HDP version.              |
| /opt/cloudera | 100 GB for CM<br>100 GB for all hosts | Minimum space required for each installed and retained CDP version. |

| Partition     | Storage  | Detail   |
|---------------|--|--|
| /usr/hdp      | 35 GB  | If you are upgrading from HDP 2.6.5 to CDP 7.1.x, there is an interim step to upgrade to HDP intermediate bits. You would need a minimum 30 GB available space to make the transition.   |
| /var/log      | 200 GB - 500 GB  | Minimum space required for storing the logs.<br> <b>Note:</b> Ensure this is not part of the root OS partition.   |
| /var/*        | 2 GB   | Minimum space required for Cloudera Manager agent and for the Cloudera components.   |
| /tmp          | 20GB   | At least 20 GB of free space required for storing the temp data by CLDR services.  |
| /<data-dir>   | Varies   | Datanode, Kafka Logs, Namenode Image and Edits, Journal Node, ZooKeeper.<br><br>These should all be on separate mounts to avoid disk issues with the os partition.<br><br>Performance Impact: For the Namenode, Journal Node, and ZooKeeper data directories, these should be on dedicated disks (and mounts) for the most optimal performance of these critical services. Disk contention with other write operations will have an impact on these services.<br><br>Performance Impact: Data-directories for Datanode and Kafka-Logs should be simple JBOD drives. RAID support has a support impact and these services are storage redundant at the service level. RAID is NOT recommended for these service data directories. |
| /<yarn-local> | Varies by workload<br>“yarn.nodemanager.local-dirs” Between 5-25% of host storage depending on workload types. | Applications that are heavy in MR technologies will benefit tremendously by using the “SSD” storage.   |

For more information on the hardware requirements for Runtime components, see [Cloudera Runtime](#).

For more information on the CM server storage requirements, see [Cloudera Manager Server](#).



For more information on the CDP Private Cloud Base requirements, see [CDP Private Cloud Base Requirements and Supported Versions](#).

## Downloading and Publishing the Package Repository

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

### Cloudera Manager 7.7.1

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

To download the files for a Cloudera Manager release, download the repository tarball for your operating system. Then unpack the tarball, move the files to the web server directory, and modify file permissions. For example:

- 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/cm7/7.7.1/repo-as-tarball/cm7.7.1-redhat7.tar.gz
```

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

```
$ tar xvfz cm7.7.1-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. If you do not see anything, your Web server may have been configured to not show indexes.

## Cloudera Manager 7.6.7

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

To download the files for a Cloudera Manager release, download the repository tarball for your operating system. Then unpack the tarball, move the files to the web server directory, and modify file permissions. For example:

- 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/cm7/7.6.7/repo-as-tarball/cm7.6.7-redhat7.tar.gz
```

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

```
$ tar xvfz cm7.6.7-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. If you do not see anything, your Web server may have been configured to not show indexes.

## Cloudera Manager 7.6.1

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

To download the files for a Cloudera Manager release, download the repository tarball for your operating system. Then unpack the tarball, move the files to the web server directory, and modify file permissions. For example:

- 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/cm7/7.6.1/repo-as-tarball/cm7.6.1-redhat7.tar.gz
```

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

```
$ tar xvfz cm7.6.1-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. If you do not see anything, your Web server may have been configured to not show indexes.

### Cloudera Manager 7.4.4

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

To download the files for a Cloudera Manager release, download the repository tarball for your operating system. Then unpack the tarball, move the files to the web server directory, and modify file permissions. For example:

- 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/cm7/7.4.4-24429768/repo-as-tarball/cm7.4.4-redhat7.tar.gz
```

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

```
$ tar xvfz cm7.4.4-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/](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.

### Cloudera Manager 7.3.1

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

To download the files for a Cloudera Manager release, download the repository tarball for your operating system. Then unpack the tarball, move the files to the web server directory, and modify file permissions. For example:

- 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/cm7/7.3.1/repo-as-tarball/cm7.3.1-redhat7.tar.gz
```

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

```
$ tar xvfz cm7.3.1-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/](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.

## Cloudera Manager 7.2.4\_patch-4455



**Note:** This patch version of Cloudera Manager is applicable only if you are upgrading from HDP 2.6.5 to CDP 7.1.4.

### 1. Cloudera Manager 7.2.4\_patch-4455

Download the patch from

```
http://[***username***]:[***password***]@bits.cloudera.com/ab44c04a/patch-4455/
```

```
tar xvfz cm 7.2.4_patch4455-redhat7.tar.gz -C /var/www/html/cloudera-repos/cm7 --strip-components=1
```

```
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. If you do not see anything, your Web server may have been configured to not show indexes.



**Important:** Include the `allkeys.asc` file at the top level of the package repository. The `allkeys.asc` file is included in the `repo-as-tarball` file. Ensure to include `allkeys.asc` file, if you are manually copying the package files between hosts.

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

## Downloading and Publishing the Parcel Repository

Download the parcels that you want to install and publish the parcel directory.

### Procedure

1. Download `manifest.json` and the parcel files for the product you want to install:

#### Runtime 7.1.7 SP2 - 7.1.7.2000

Apache Impala, Apache Kudu, Apache Spark 2, and Cloudera Search are included in the Runtime parcel. 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/cdh7/7.1.7.2000/parcels/ -P /var/www/html/cloudera-repos
sudo chmod -R ugo+rX /var/www/html/cloudera-repos/p/cdh7
```

#### Runtime 7.1.7 SP1 - Use this if you are upgrading from CDP 7.1.7 to CDP 7.1.7 SP1

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/cdh7/7.1.7.1000/parcels/ -P /var/www/html/cloudera-repos
```

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

#### Runtime 7.1.7.78 - Use this if you are upgrading from HDP 2.6.5 to CDP 7.1.7

Apache Impala, Apache Kudu, Apache Spark 2, and Cloudera Search are included in the Runtime parcel. 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/cdh7/7.1.7.78/parcels/ -P /var/www/html/cloudera-repos
sudo chmod -R ugo+rX /var/www/html/cloudera-repos/p/cdh7
```

#### Runtime 7.1.6.0 - Use this if you are upgrading from HDP 2.6.5 to CDP 7.1.6

Apache Impala, Apache Kudu, Apache Spark 2, and Cloudera Search are included in the Runtime parcel. 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/cdh7/7.1.6.0/parcels/ -P /var/www/html/cloudera-repos
sudo chmod -R ugo+rX /var/www/html/cloudera-repos/p/cdh7
```

#### Runtime 7.1.4.6 - Use this if you are upgrading from HDP 2.6.5 to 7.1.4.6

Apache Impala, Apache Kudu, Apache Spark 2, and Cloudera Search are included in the Runtime parcel. 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/cdh7/7.1.4.6/parcels/ -P /var/www/html/cloudera-repos
sudo chmod -R ugo+rX /var/www/html/cloudera-repos/p/cdh7
```

### Sqoop Connectors

To download the parcels for a Sqoop Connector release, run the following commands on the Web server host. This example uses the latest available Sqoop Connectors:

```
sudo mkdir -p /var/www/html/cloudera-repos
sudo wget --recursive --no-parent --no-host-directories http://archive.cloudera.com/sqoop-connectors/parcels/latest/ -P /var/www/html/cloudera-repos
sudo chmod -R ugo+rX /var/www/html/cloudera-repos/sqoop-connectors
```

If you want to create a repository for a different Sqoop Connector release, replace latest with the Sqoop Connector version that you want. You can see a list of versions in the parcels parent directory.

2. 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.

## Hadoop Users (user:group) and Kerberos Principals

During the Cloudera Manager installation process, several Linux user accounts and groups are created by default. These are listed in the table below. Integrating the cluster to use Kerberos for authentication requires creating Kerberos principals and keytabs for these user accounts.

**Table 2: Users and Groups**

| Component (Version)        | Unix User ID | Groups        | Functionality   |
|----------------------------|--------------|---------------|---|
| Apache Atlas               | atlas        | atlas, hadoop | Apache Atlas by default has atlas as user and group. It is configurable   |
| Apache Flink               | flink        | flink         | The Flink Dashboard runs as this user.  |
| Apache HBase               | hbase        | hbase         | The Master and the RegionServer processes run as this user.   |
| Apache HBase Indexer       | hbase        | hbase         | The indexer servers are run as this user.   |
| Apache HDFS                | hdfs         | hdfs, hadoop  | The NameNode and DataNodes run as this user, and the HDFS root directory as well as the directories used for edit logs should be owned by it.   |
| Apache Hive<br>Hive on Tez | hive         | hive          | The HiveServer2 process and the Hive Metastore processes run as this user. A user must be defined for Hive access to its Metastore DB (for example, MySQL or Postgres) but it can be any identifier and does not correspond to a Unix uid. This is <code>javax.jdo.option.ConnectionUserName</code> in <code>hive-site.xml</code> . |
| Apache Impala              | impala       | impala, hive  | Impala services run as this user.   |
| Apache Kafka               | kafka        | kafka         | Kafka brokers, mirrorMaker, and Connect workers run as this user.   |
| Apache Knox                | knox         | knox          | Apache Knox Gateway Server runs as this user  |
| Apache Kudu                | kudu         | kudu          | Kudu services run as this user.   |
| Apache Livy                | livy         | livy          | The Livy Server process runs as this user   |
| Apache NiFi                | nifi         | nifi          | Runs as the nifi user   |
| Apache NiFi Registry       | nifiregistry | nifiregistry  | Runs as the nifiregistry user   |
| Apache Oozie               | oozie        | oozie         | The Oozie service runs as this user.  |
| Apache Ozone               | hdfs         | hdfs, hadoop  | Ozone Manager, Storage Container Manager (SCM), Recon and Ozone Datanodes run as this user.   |
| Apache Parquet             | ~            | ~             | No special users.   |
| Apache Phoenix             | phoenix      | phoenix       | The Phoenix Query Server runs as this user  |

| Component (Version)             | Unix User ID   | Groups         | Functionality  |
|---------------------------------|----------------|----------------|--|
| Apache Ranger                   | ranger         | ranger, hadoop | Ranger Admin, Usersync and Tagsync services by default have ranger as user and ranger, hadoop as groups. It is configurable.   |
| Apache Ranger KMS               | kms            | kms            | Ranger KMS runs with kms user and group. It is configurable.   |
| Apache Ranger Raz               | rangerraz      | ranger         | Ranger Raz runs with rangerraz user and is part of the ranger group.   |
| Apache Ranger RMS               | rangerrms      | ranger         | Ranger RMS runs with rangerrms user and is part of the ranger group.   |
| Apache Solr                     | solr           | solr           | The Solr processes run as this user.   |
| Apache Spark                    | spark          | spark          | The Spark History Server process runs as this user.  |
| Apache Sqoop                    | sqoop          | sqoop          | This user is only for the Sqoop1 Metastore, a configuration option that is not recommended.  |
| Apache YARN                     | yarn           | yarn, hadoop   | Without Kerberos, all YARN services and applications run as this user. The LinuxContainerExecutor binary is owned by this user for Kerberos.   |
| Apache Zeppelin                 | zeppelin       | zeppelin       | The Zeppelin Server process runs as this user  |
| Apache ZooKeeper                | zookeeper      | zookeeper      | The ZooKeeper processes run as this user. It is not configurable.  |
| Cloudera Manager (all versions) | cloudera-scm   | cloudera-scm   | Clusters managed by Cloudera Manager run Cloudera Manager Server, monitoring roles, and other Cloudera Server processes as cloudera-scm. Requires keytab file named cmf.keytab because name is hard-coded in Cloudera Manager. |
| Cruise Control                  | cruisecontrol  | hadoop         | The Cruise Control process runs as this user.  |
| HttpFS                          | httpfs         | httpfs         | The HttpFS service runs as this user. See "HttpFS authentication" for instructions on how to generate the merged httpfs-http.keytab file.  |
| Hue                             | hue            | hue            | Hue services run as this user.   |
| Hue Load Balancer               | apache         | apache         | The Hue Load balancer has a dependency on the apache2 package that uses the apache user name. Cloudera Manager does not run processes using this user ID.  |
| Key Trustee Server              | keytrustee     | keytrustee     | The Key Trustee Server service runs as this user.  |
| Schema Registry                 | schemaregistry | hadoop         | The Schema Registry process runs as this user.   |
| Streams Messaging Manager       | streamsmgmr    | streamsmgmr    | The Streams Messaging Manager processes runs as this user.   |

| Component (Version)         | Unix User ID  | Groups        | Functionality  |
|-----------------------------|---------------|---------------|--|
| Streams Replication Manager | streamsrepmgr | streamsrepmgr | The Streams Replication Manager processes runs as this user. |

### Keytabs and Keytab File Permissions

Linux user accounts, such as hdfs, are mapped to the username portion of the Kerberos principal names, as follows:

```
username/host.example.com@EXAMPLE.COM
```

For example, the Kerberos principal for Apache Hive would be:

```
hive/host.example.com@EXAMPLE.COM
```

Keytabs that contain multiple principals are merged automatically from individual keytabs by Cloudera Manager. If you override a service configuration to not use the CM-provided keytab, then you must ensure that all the principals required for the given role instance on a specific host are merged together in the keytab file you deploy manually on that host.

For example, for Filename (\*.keytab), the Atlas keytab filename would be atlas.keytab, HBase would be hbase.keytab, and Cloudera Manager would be cmf.keytab and scm.keytab.

Keytab File Owner:Group matters when Cloudera Manager starts a role. For example, Cloudera Manager starts the role "DataNode". Cloudera Manager launches the DataNode process as a user (here, "hdfs"). Because that process needs to access the HDFS keytab, Cloudera Manager puts the HDFS keytab in the DataNode's process directory, and the keytab is given the owner:group that is listed in the table. Thus, the DataNode process properly owns the keytab file.

The tables below lists the usernames to use for Kerberos principal names, for clusters managed by Cloudera Manager.

#### Apache Atlas

**Role:** atlas-ATLAS\_SERVER

**Kerberos Principals**

atlas

**Filename (\*.keytab)**

atlas

**Keytab File Owner:Group**

atlas:atlas

**File Permission (octal)**

600

#### Apache Flink

**Role:** flink

**Kerberos Principals**

flink

**Filename (\*.keytab)**

flink

**Keytab File Owner:Group**

flink:flink

**File Permission (octal)**

600



**Apache HBase****Role: hbase-HBASETHRIFTSERVER****Kerberos Principals**

hbase, HTTP

**Filename (\*.keytab)**

hbase, HTTP

**Keytab File Owner:Group**

hbase:hbase

**File Permission (octal)**

600

**Role: hbase-REGIONSERVER****Kerberos Principals**

hbase, HTTP

**Filename (\*.keytab)**

hbase, HTTP

**Keytab File Owner:Group**

hbase:hbase

**File Permission (octal)**

600

**Role: hbase-HBASERESTSERVER****Kerberos Principals**

hbase, HTTP

**Filename (\*.keytab)**

hbase, HTTP

**Keytab File Owner:Group**

hbase:hbase

**File Permission (octal)**

600

**Role: hbase-MASTER****Kerberos Principals**

hbase, HTTP

**Filename (\*.keytab)**

hbase, HTTP

**Keytab File Owner:Group**

hbase:hbase

**File Permission (octal)**

600

**Apache HBase indexer****Role: ks\_indexer-HBASE\_INDEXER****Kerberos Principals**

hbase, HTTP

**Filename (\*.keytab)**

hbase

**Keytab File Owner:Group**

hbase:hbase

**File Permission (octal)**

600

**Apache HDFS****Role: hdfs-NAMENODE****Kerberos Principals**

hdfs, HTTP

**Filename (\*.keytab)**

hdfs

**Keytab File Owner:Group**

hdfs:hdfs

**File Permission (octal)**

600

**Role: hdfs-DATANODE****Kerberos Principals**

hdfs, HTTP

**Filename (\*.keytab)**

hdfs

**Keytab File Owner:Group**

hdfs:hdfs

**File Permission (octal)**

600

**Role: hdfs-SECONDARYNAMENODE****Kerberos Principals**

hdfs, HTTP

**Filename (\*.keytab)**

hdfs

**Keytab File Owner:Group**

hdfs:hdfs

**File Permission (octal)**

600

**Apache Hive, Hive on Tez****Role: hive-HIVESERVER2****Kerberos Principals**

hive

**Filename (\*.keytab)**

hive

**Keytab File Owner:Group**

hive:hive

**File Permission (octal)**

600

**Role: hive-HIVEMETASTORE****Kerberos Principals**

hive

**Filename (\*.keytab)**

hive

**Keytab File Owner:Group**

cloudera-scm:cloudera-scm

**File Permission (octal)**

600

**Apache Impala****Role: impala-STATESTORE****Kerberos Principals**

impala, HTTP

**Filename (\*.keytab)**

impala

**Keytab File Owner:Group**

impala:impala

**File Permission (octal)**

600

**Role: impala-CATALOGSERVER****Kerberos Principals**

impala, HTTP

**Filename (\*.keytab)**

impala

**Keytab File Owner:Group**

impala:impala

**File Permission (octal)**

600

**Role: impala-IMPALAD****Kerberos Principals**

impala, HTTP

**Filename (\*.keytab)**

impala

**Keytab File Owner:Group**

impala:impala

**File Permission (octal)**

600

**Apache Kafka****Role: kafka-KAFKA\_BROKER****Kerberos Principals**

kafka

**Filename (\*.keytab)**

kafka

**Keytab File Owner:Group**

kafka:kafka

**File Permission (octal)**

600

**Role: kafka-KAFKA\_MIRROR\_MAKER****Kerberos Principals**

kafka\_mirror\_maker

**Filename (\*.keytab)**

kafka

**Keytab File Owner:Group**

kafka:kafka

**File Permission (octal)**

600

**Role: kafka-KAFKA\_CONNECT****Kerberos Principals**

kafka

**Filename (\*.keytab)**

kafka

**Keytab File Owner:Group**

kafka:kafka

**File Permission (octal)**

600

**Apache Knox****Role: knox-KNOX\_GATEWAY****Kerberos Principals**

knox, HTTP

**Filename (\*.keytab)**

hbase

**Keytab File Owner:Group**

knox:knox

**File Permission (octal)**

600

**Apache Kudu****Role: kudu-KUDU\_MASTER****Kerberos Principals**

kudu

**Filename (\*.keytab)**

kudu

**Keytab File Owner:Group**

kudu:kudu

**File Permission (octal)**

600

**Role: kudu-KUDU\_TSERVER****Kerberos Principals**

kudu

**Filename (\*.keytab)**

kudu

**Keytab File Owner:Group**

kudu:kudu

**File Permission (octal)**

600

**Apache Livy****Role: livy-LIVY\_SERVER****Kerberos Principals**

livy

**Filename (\*.keytab)**

livy

**Keytab File Owner:Group**

livy:livy

**File Permission (octal)**

600

**Apache NiFi****Role: nifi****Kerberos Principals**

nifi, HTTP

**Filename (\*.keytab)**

nifi

**Keytab File Owner:Group**

nifi:nifi

**File Permission (octal)**

600

**Apache NiFi Registry****Role: nifiregistry****Kerberos Principals**

nifiregistry, HTTP

**Filename (\*.keytab)**

nifiregistry

**Keytab File Owner:Group**

nifiregistry:nifiregistry

**File Permission (octal)**

600

**Apache Oozie****Role: oozie-OOZIE\_SERVER****Kerberos Principals**

oozie, HTTP

**Filename (\*.keytab)**

oozie

**Keytab File Owner:Group**

oozie:oozie

**File Permission (octal)**

600

**Apache Ozone**

**Role: ozone-OZONE\_MANAGER**

**Kerberos Principals**

om, HTTP

**Filename (\*.keytab)**

ozone

**Keytab File Owner:Group**

hdfs:hdfs

**File Permission (octal)**

600

**Role: ozone-STORAGE\_CONTAINER\_MANAGER**

**Kerberos Principals**

scm, HTTP

**Filename (\*.keytab)**

ozone

**Keytab File Owner:Group**

hdfs:hdfs

**File Permission (octal)**

600

**Role: ozone-OZONE\_DATANODE**

**Kerberos Principals**

dn, HTTP

**Filename (\*.keytab)**

ozone

**Keytab File Owner:Group**

hdfs:hdfs

**File Permission (octal)**

600

**Role: ozone-OZONE\_RECON**

**Kerberos Principals**

recon, HTTP

**Filename (\*.keytab)**

ozone

**Keytab File Owner:Group**

hdfs:hdfs

**File Permission (octal)**

600

**Role: ozone-S3\_GATEWAY**

**Kerberos Principals**

HTTP

**Filename (\*.keytab)**

ozone

**Keytab File Owner:Group**

hdfs:hdfs

**File Permission (octal)**

600

### Apache Phoenix

**Role: phoenix-PHOENIX\_QUERY\_SERVER**

**Kerberos Principals**

phoenix, HTTP

**Filename (\*.keytab)**

phoenix

**Keytab File Owner:Group**

phoenix:phoenix

**File Permission (octal)**

600

### Apache Ranger

**Role: ranger-RANGER\_ADMIN**

**Kerberos Principals**

rangeradmin, rangerlookup, HTTP

**Filename (\*.keytab)**

ranger

**Keytab File Owner:Group**

ranger:ranger

**File Permission (octal)**

600

**Role: ranger-RANGER\_USERSYNC**

**Kerberos Principals**

rangerusersync

**Filename (\*.keytab)**

ranger

**Keytab File Owner:Group**

ranger:ranger

**File Permission (octal)**

600

**Role: ranger-RANGER\_TAGSYNC**

**Kerberos Principals**

rangertagsync

**Filename (\*.keytab)**

ranger

**Keytab File Owner:Group**

ranger:ranger

**File Permission (octal)**

600

**Apache Ranger KMS**

**Role:** ranger-RANGER\_TAGSYNC

**Kerberos Principals**

rangerkms, HTTP

**Filename (\*.keytab)**

ranger\_kms

**Keytab File Owner:Group**

kms:kms

**File Permission (octal)**

600

**Apache Ranger Raz**

**Role:** ranger-RANGER\_RAZ

**Kerberos Principals**

rangerraz, HTTP

**Filename (\*.keytab)**

rangerraz

**Keytab File Owner:Group**

ranger:rangerraz

**File Permission (octal)**

600

**Apache Ranger RMS**

**Role:** ranger-RANGER\_RMS

**Kerberos Principals**

rangerrms

**Filename (\*.keytab)**

rangerrms

**Keytab File Owner:Group**

ranger:rangerrms

**File Permission (octal)**

600

**Apache Solr**

**Role:** solr-SOLR\_SERVER

**Kerberos Principals**

solr, HTTP

**Filename (\*.keytab)**

solr

**Keytab File Owner:Group**



solr:solr

**File Permission (octal)**

600

**Apache Spark**

**Role: spark\_on\_yarn-SPARK\_YARN\_HISTORY\_SERVER**

**Kerberos Principals**

spark

**Filename (\*.keytab)**

spark

**Keytab File Owner:Group**

spark:spark

**File Permission (octal)**

600

**Apache YARN**

**Role: yarn-NODEMANAGER**

**Kerberos Principals**

yarn, HTTP

**Filename (\*.keytab)**

yarn

**Keytab File Owner:Group**

yarn:hadoop

**File Permission (octal)**

644

**Role: yarn-RESOURCEMANAGER**

**Kerberos Principals**

yarn, HTTP

**Filename (\*.keytab)**

yarn

**Keytab File Owner:Group**

yarn:hadoop

**File Permission (octal)**

600

**Role: yarn-JOBHISTORY**

**Kerberos Principals**

mapred

**Filename (\*.keytab)**

mapred

**Keytab File Owner:Group**

yarn:hadoop

**File Permission (octal)**

600

**Apache Zeppelin****Role: zeppelin-ZEPPELIN\_SERVER****Kerberos Principals**

zeppelin, HTTP

**Filename (\*.keytab)**

zeppelin

**Keytab File Owner:Group**

zeppelin:zeppelin

**File Permission (octal)**

600

**Apache ZooKeeper****Role: zookeeper-server****Kerberos Principals**

zookeeper

**Filename (\*.keytab)**

zookeeper

**Keytab File Owner:Group**

zookeeper:zookeeper

**File Permission (octal)**

600

**Cloudera Management****Role: cloudera-mgmt-REPORTSMANAGER****Kerberos Principals**

hdfs

**Filename (\*.keytab)**

headlamp

**Keytab File Owner:Group**

cloudera-scm:cloudera-scm

**File Permission (octal)**

600

**Role: cloudera-mgmt-SERVICEMONITOR****Kerberos Principals**

hue

**Filename (\*.keytab)**

cmon

**Keytab File Owner:Group**

cloudera-scm:cloudera-scm

**File Permission (octal)**

600

**Role: cloudera-mgmt-ACTIVITYMONITOR****Kerberos Principals**

hue

**Filename (\*.keytab)**

cmon

**Keytab File Owner:Group**

cloudera-scm:cloudera-scm

**File Permission (octal)**

600

**Cloudera Manager**

**Kerberos Principals**

cloudera-scm, HTTP

**Filename (\*.keytab)**

cmf, scm

**Keytab File Owner:Group**

cloudera-scm:cloudera-scm

**File Permission (octal)**

600

**CruiseControl**

**Role: cruise\_control-CRUISE\_CONTROL\_SERVER**

**Kerberos Principals**

cruisecontrol, kafka, HTTP

**Filename (\*.keytab)**

cruise\_control

**Keytab File Owner:Group**

cruisecontrol:hadoop

**File Permission (octal)**

600

**HttpFS**

**Role: hdfs-HTTPFS**

**Kerberos Principals**

httpfs, HTTP

**Filename (\*.keytab)**

httpfs

**Keytab File Owner:Group**

httpfs:httpfs

**File Permission (octal)**

600

**Hue**

**Role: hue-KT\_RENEWER**

**Kerberos Principals**

hue

**Filename (\*.keytab)**

hue

**Keytab File Owner:Group**

hue:hue

**File Permission (octal)**

600

**Schema Registry**

**Role:** schemaregistry-SCHEMA\_REGISTRY\_SERVER

**Kerberos Principals**

schemaregistry, HTTP

**Filename (\*.keytab)**

schemaregistry

**Keytab File Owner:Group**

schemaregistry:hadoop

**File Permission (octal)**

600

**Streams Messaging Manager**

**Role:** streams\_messaging\_manager-STREAMS\_MESSAGING\_MANAGER\_SERVER

**Kerberos Principals**

streamsmgmgr, HTTP

**Filename (\*.keytab)**

streams\_messaging\_manager

**Keytab File Owner:Group**

streamsmgmgr:streamsmgmgr

**File Permission (octal)**

600

**Streams Replication Manager**

**Role:** streams\_replication\_manager-STREAMS\_REPLICATION\_MANAGER\_DRIVER

**Kerberos Principals**

streamsrepmgr

**Filename (\*.keytab)**

streams\_replication\_manager

**Keytab File Owner:Group**

streamsrepmgr:streamsrepmgr

**File Permission (octal)**

600

**Role:** streams\_replication\_manager-STREAMS\_REPLICATION\_MANAGER\_SERVICE

**Kerberos Principals**

streamsrepmgr

**Filename (\*.keytab)**

streams\_replication\_manager

**Keytab File Owner:Group**

streamsrepmgr:streamsrepmgr

**File Permission (octal)**

600

## Create Service Principals and Keytab Files for HDP

This section is optional. During the HDP 2.6.5 to HDP intermediate bits upgrade, Ambari can generate the principals and keytabs. However, before upgrading, you can manually generate the principals and keytabs. First, create the principal using mandatory naming conventions and then create the keytab file with the principal's information. Lastly, copy the keytab file to the keytab directory on the appropriate service host.

To create a service principal, use the kadmin utility. The kadmin utility is a command-line driven utility where you can run Kerberos commands to manipulate the central database. To start kadmin, run the following commands:

1. `'kadmin $USER/admin@REALM'`
2. `kadmin: addprinc -randkey $principal_name/$service-host-FQDN@$hadoop.realm`



### Note:

- a. You must have a principal with administrative permissions to run the above commands.
- b. The randkey is used to generate the password.
- c. The `$principal_name` part of the name must match the values in the table below.

In the example mentioned in step 2, each service principal's name is appended with a fully qualified domain name of the host on which the principal is running. This is to provide a unique principal name for services that run on multiple hosts, like DataNodes and TaskTrackers. The addition of the hostname serves to distinguish, for example, a request from DataNode A from a request from DataNode B. This is important for two reasons:

- If the Kerberos credentials for one DataNode are compromised, it does not automatically compromise all other DataNodes.
- If multiple DataNodes have the same principal and are simultaneously connecting to the NameNode, and if the Kerberos authenticator sent has the same timestamp, then the authentication is rejected as a replay request.



**Note:** The NameNode, Secondary NameNode, and Oozie require two principals each.

If you are configuring High Availability (HA) for a Quorum-based NameNode, you must also generate a principle (`jn/$FQDN`) and keytab (`jn.service.keytab`) for each JournalNode. JournalNode also requires the keytab for its HTTP service. If the JournalNode is deployed on the same host as a NameNode, the same keytab file (`spnego.service.keytab`) can be used for both. In addition, HA requires two NameNodes. Both the active and standby NameNodes require their own principal and keytab files. The service principals of the two NameNodes can share the same name, specified with the `dfs.namenode.kerberos.principal` property in `hdfs-site.xml`, but the NameNodes still have different fully qualified domain names.

| Service   | Component/Role          | Principal Name  | Mandatory Keytab Filename |
|-----------|-------------------------|-----------------|---------------------------|
| HDFS      | NameNode                | nn/\$FQDN       | nn.service.keytab         |
|           | SecondaryNameNode       | nn/\$FQDN       | nn.service.keytab         |
|           | DataNode                | dn/\$FQDN       | dn.service.keytab         |
|           | Journal Server*         | jn/\$FQDN       | jn.service.keytab         |
|           | NameNode HTTP           | HTTP/\$FQDN     | spnego.service.keytab     |
|           | SecondaryNameNode HTTP  | HTTP/\$FQDN     | spnego.service.keytab     |
| MapReduce | MR2 History Server      | jhs/\$FQDN      | nm.service.keytab         |
|           | MR2 History Server HTTP | HTTP/\$FQDN     | spnego.service.keytab     |
| YARN      | Node Manager            | nm/\$FQDN       | nm.service.keytab         |
|           | Resource Manager        | rm/\$FQDN       | rm.service.keytab         |
|           | YARN Timeline Server    | yarn-ats/\$FQDN | yarn-ats.service.keytab   |
|           | HTTP                    | HTTP/\$FQDN     | spnego.service.keytab     |
| Oozie     | Oozie Server            | oozie/\$FQDN    | oozie..service.keytab     |

| Service    | Component/Role      | Principal Name        | Mandatory Keytab Filename        |
|------------|---------------------|-----------------------|----------------------------------|
|            | Oozie HTTP          | HTTP/\$FQDN           | spnego.service.keytab            |
| Hive       | HiverServer2, HMS   | hive/\$FQDN           | hive.service.keytab              |
|            | Hive HTTP           | HTTP/\$FQDN           | spnego.service.keytab            |
| HBase      | HBase Master Server | hbase/\$FQDN          | hbase.service.keytab             |
|            | HBase RegionServer  | hbase/\$FQDN          | hbase.service.keytab             |
| Kafka      | Kafka Broker        | kafka/\$FQDN          | kafka.service.keytab             |
| Zeppelin   | Zeppelin Server     | zeppelin/\$FQDN       | zeppelin.service.keytab          |
| Zookeeper  |                     | zookeeper/\$FQDN      | zk.service.keytab                |
| Knox       |                     | knox/\$FQDN           | knox.service.keytab              |
| Ranger     | Admin Server        | rangeradmin/\$FQDN    | rangeradmin.service.keytab       |
|            | Lookup Server       | rangerlookup/\$FQDN   | rangerlookup.service.keytab      |
|            | KMS                 | rangerkms/\$FQDN      | rangerkms.service.keytab         |
|            | UserSync            | rangerusersync/\$FQDN | rangerusersync.service.keytab    |
|            | TagSync             | rangertagsync/\$FQDN  | rangertagsync.service.keytab     |
| AMS        |                     | amshbase/\$FQDN       | ams-hbase.master.keytab          |
|            |                     | amsmon/\$FQDN         | ams.collector.keytab             |
|            |                     | amszk/\$FQDN          | ams-zk.service.keytab            |
| Spark2     |                     | spark/\$FQDN          | spark.service.keytab             |
| Druid      |                     | druid/\$FQDN          | druid.service.keytab             |
| Infra-Solr |                     | infra-solr/\$FQDN     | ambari-infra-solr.service.keytab |
| Atlas      |                     | atlas/\$FQDN          | atlas.service.keytab             |
| Livy       |                     | livy/\$FQDN           | livy.service.keytab              |

\* Only required if you are setting up NameNode HA. For example, to create the principal for a DataNode service, run the command `kadmin: addprinc -randkey dn/$datanode-host@$hadoop.realm`

3. Extract the related keytab file and place it in the keytab directory of the respective components. The default directory is `/etc/krb5.keytab`.

- `kadmin: xst -k $keytab_file_name $principal_name/fully.qualified.domain.name`


You must use the mandatory names for the `$keytab_file_name` variable shown in the table above. For example, to create the keytab files for the NameNode, run the command `kadmin: xst -k nn.service.keytab nn/$namenode-host` `kadmin: xst -k spnego.service.keytab HTTP/$namenode-host`

After creating the keytab files, copy the keytab files to the keytab directory of the respective service hosts.

4. On each service in your cluster, verify that the correct keytab files and principals are associated with the correct service using the `klist` command. For example, on the NameNode, run the command `klist -k -t /etc/security/nn.service.keytab`

## Sample data ingestion


Cloudera recommends you use a subset of your workload or any sample data for any jobs or queries and do a benchmarking. You can record the time taken for the critical jobs and compare the performance of pre and post upgrade setup.



Cloudera recommends you use a subset of your workload or any sample data for any jobs or queries and do a benchmarking if required. You can record the time taken for the critical jobs and compare the performance of pre and post upgrade setup.

## Expediting the Hive upgrade

Preparing the Hive metastore for the upgrade can take a long time. Checking and correcting your Hive metastore partitions and SERDE definitions is critical for a successful upgrade. If you have many tables and partitions, it might be difficult to manually identify these problems. The free Hive Upgrade Check tool helps identify these problems.



The Hive Upgrade Check tool is Community software that scans your Hive metastore to identify potential upgrade problems. You can also use the tool to perform the following tasks:

- Convert legacy managed tables (non-acid) to external tables.
- Report potential problems, such as tables that do not have matching HDFS directories, to resolve before the upgrade.

The cluster upgrade to CDP runs the Hive Strict Managed Migration (HSMM) process that performs the same tasks. During the cluster upgrade, you can skip the HSMM process, migrating none of your tables and database definitions.

## Overview of the expedited Hive upgrade

You perform tasks before and after the Hive migration to hasten the upgrade. The sequence of steps involved in expediting the Hive upgrade includes identifying problems in tables and databases before upgrading, modifying the

HSMM process to prevent migration of your tables and databases, and completing the cluster upgrade. After the upgrade to CDP, you migrate the tables and databases.

1. Prepare tables for migration, identifying potential migration problems using the Hive Upgrade Check tool.
2. Decide to expedite the upgrade by not migrating your databases and tables during the upgrade.
3. Modify the HSMM process to prevent migration of your tables and databases during the upgrade.
4. Start the upgrade wizard to upgrade your cluster.

If you did not migrate all tables, do so after the upgrade to CDP as follows:

1. Prepare your databases and tables for migration using the Hive Upgrade Check tool.
2. Migrate tables and databases to CDP.

You cannot use unmigrated tables in CDP.

## Handle Missing Table or Partition Locations

You need to identify missing table or partition locations, or both, to prevent upgrade failure. If the table and partition locations do not exist in the file system, you must either create a replacement partition directory (recommended) or drop the table and partition.

### About this task

You perform this step if you did not modify the HSMM process to expedite the Hive upgrade.

### Procedure

Ensure the table and partition locations exist on the file system. If these locations don't exist either create a replacement partition directory (recommended) or drop the table and partition.

## Managed Table Location Mapping

A managed table location must map to one managed table only. If multiple managed tables point to the same location, upgrade problems occur.

## Check SERDE Definitions and Availability

Ensure correct Serde definitions and a reference to a SERDE exists to ensure a successful upgrade.

### About this task

You perform this step if you do not modify the HSMM process for expediting the Hive upgrade.

### Procedure

1. Check Serde definitions for correctness and check for SERDE availability.
2. Correct any problems found as follows:
  - Remove the table having the problematic SERDE.
  - Ensure the SERDE is available during the upgrade, so the table can be evaluated.

## Make Tables SparkSQL Compatible

Non-Acid, managed tables in ORC or in a Hive Native (but non-ORC) format that are owned by the POSIX user hive will not be SparkSQL-compatible after the upgrade unless you perform manual conversions.

### About this task

If your table is a managed, non-ACID table, you can convert it to an external table using this procedure (recommended). After the upgrade, you can easily convert the external table to an ACID table, and then use the Hive Warehouse Connector to access the ACID table from Spark.

Take one of the following actions.



- Convert the tables to external Hive tables before the upgrade.  
`ALTER TABLE ... SET TBLPROPERTIES('EXTERNAL'='TRUE','external.table.purge'='true')`
- Change the POSIX ownership to an owner other than hive.  
You will need to convert managed, ACID v1 tables to external tables after the upgrade.

## Understanding the Hive upgrade

You need an understanding of the Hive Strict Managed Migration (HSMM) and the Hive Upgrade Check tool for a successful upgrade.

### HSMM vs the Hive Upgrade Check tool

It is difficult to estimate how long the Hive Strict Managed Migration will take. The following factors are just a few that might affect how long it takes:

- Number of managed tables
- Core processing power
- Backend metastore database speed

The process runs across all Hive metastore databases and tables by default, identifying managed tables that need to undergo compaction or conversion to Hive 3 ACID V2 tables.

Consider expediting the upgrade process if one of the following conditions exist:

- You have few, or no, ACID tables but do have many legacy managed tables in your environment.
- Reducing downtime is critical, and justifies the extra effort to expedite the upgrade process.

### Why upgrading takes so long

The underlying Hive upgrade process [Hive Strict Managed Migration](#) (HSMM) is an Apache Hive conversion utility that makes adjustments to Hive tables under the enhanced and strict Hive 3 environment to meet the needs of the most demanding workloads and governance requirements for Data Lake implementations. There are some changes to the standard behaviors in Hive table definitions and locations. The HSMM reviews every database and table to determine if changes are needed to meet these requirements.

With systems that have been around for a while, or have adopted some ingest patterns, there may be artifacts in the metastore that cannot be reconciled, including the following artifacts:

- Tables and partitions without reciprocating storage locations
- Tables using SERDEs that have been abandoned.
- ACIDv1 tables

These tables must be fully compacted before the upgrade. If tables are not compacted, data loss is highly likely.

When these irreconcilable conditions occur, it requires manual intervention to fix problems before it can proceed.

The Hive upgrade process iterates through the databases and tables, attempting to materialize each of them using the Hive Metastore and public Thrift APIs. That creates a heavy load on the underlying metastore database and entire system.

## Modifying the HSMM to prevent migration

By default, the Hive Strict Managed Migration process (HSMM) runs across all tables and databases. But, you want to prevent HSMM from migrating tables and databases. To change the default, you must directly modify the Ambari-DC Hive Upgrade Process.

### Before you begin

You upgraded to Ambari-DC, but have not yet performed the HDP upgrade of the platform stack from Hive 1 to Hive 3.

### About this task

Perform this procedure if you want to expedite the upgrade process.

### Procedure

1. Open the following script on the Ambari Server host for editing: `/var/lib/ambari-server/resources/stacks/HDP/3.0/services/HIVE/package/scripts/post_upgrade.py`
2. Make the changes as shown in the following example:

FROM

```
cmd = format("{hive_script} --config /etc/hive/conf --service strictman
agedmigration --hiveconf hive.strict.managed.tables=true -m automatic
--modifyManagedTables --oldWarehouseRoot /apps/hive/warehouse")
```

TO

```
cmd = format("{hive_script} --config /etc/hive/conf --service strictman
agedmigration --hiveconf hive.strict.managed.tables=true -m automatic
--modifyManagedTables --oldWarehouseRoot /apps/hive/warehouse --tableRegex
zzzzz")
```

3. Restart the Ambari Server and agents after this change but before starting the upgrade process.




**Important:** After upgrading to CDP, you must run the Hive Upgrade Check tool and do any suggested clean up on the Hive metastore, and then migrate the corrected tables to CDP.

### Related Information

[Running the Hive Upgrade Check tool](#)

## Ambari and HDP Upgrade Checklist

Before upgrading Ambari and HDP, you must review the checklist for activities that you must perform to ensure the cluster is ready for the upgrade.



## Ambari upgrade checklist

When preparing to upgrade Ambari and the HDP Cluster, Cloudera recommends that you review this checklist to confirm that your cluster operation is healthy. Attempting to upgrade a cluster that is operating in an unhealthy state can produce unexpected results.

### Pre-upgrade checklist

When preparing to upgrade Ambari and the HDP Cluster, Cloudera recommends that you review this checklist to confirm that your cluster operation is healthy. Attempting to upgrade a cluster that is operating in an unhealthy state can produce unexpected results.

Important: Always ensure that you are using the most recent version of Ambari to perform your upgrade.

- Ensure all services in the cluster are running.
- Run each Service Check (found under the Service Actions menu) and confirm they execute successfully.
- Clear all alerts, or understand why they are being generated. Remediate as necessary.
- Confirm that the start and stop options for all services are executing successfully.
- Record the time it takes to start and stop each service. The time to start and stop services is a big contributor to overall upgrade time so having this information handy is useful.
- Download the software packages prior to the upgrade. Place them in a local repository and consider using a storage proxy because multi-gigabyte downloads are required on all nodes in the cluster.
- Ensure point-in-time backups are taken of all databases that support the cluster. This includes (among others) Ambari, Hive, Ranger, Druid, Superset, and Oozie.



**Note:** To find all the download links in a centralised location, see [Software download matrix](#)

## Download cluster blueprints without hosts

Ensure you download the cluster blueprints before you begin the upgrade. The blueprint of your old cluster is used as the template of the new cluster after you transition to Cloudera Manager.

### About this task

### Procedure

To download the cluster blueprints without hosts, do the following:

```
curl ${securecurl} -H "X-Requested-By: ambari" -X GET -u ${ambariuser}:${ambaripwd} ${ambariprotocol}://${ambariserver}:${ambariport}/api/v1/clusters/${clustername}?format=blueprint > "${backupdir}/${clustername}_blueprint_without_hosts_"$(date +%Y%m%d%H%M%S)".json
```

where,

securecurl = -k for https

securecurl = -k for https

ambariuser = Ambari User Name

ambaripwd = Ambari Password

ambariprotocol = http or https

ambariserver = Ambari Server Name

ambariport = Ambari Port

clustername = The Name of the cluster

backupdir = The folder to download the blueprint

## HDP upgrade checklist

You must ensure that your HDP cluster satisfies specific prerequisites before initiating the upgrade from Ambari 2.6.2.x to Ambari 7.1.x.x.

- You must review the [Behavioral Changes](#).
- You must have root, administrative, or root-equivalent authorization on the Ambari Server host and all Ambari Agent hosts in the cluster.
- You must make a safe copy of the Ambari Server configuration file found at `/etc/ambari-server/conf/ambari.properties`.
- If your cluster is SSO-enabled, do not stop Knox before upgrading Ambari.
- You must disable the Queue Priority Support feature of YARN.

## Checklist for large clusters

If you have a large cluster with a large data volume, you should configure the Ambari NameNode restart timeout parameter on the Ambari Server host, to ensure that the Ambari requests for starting the NameNode do not timeout during an upgrade.

In a large cluster, NameNode startup processes can take a long time. NameNode startup time depends not only on host properties, but also on data volume and network parameters. To ensure that the Ambari requests for starting the NameNode do not timeout during an upgrade, you should configure the Ambari NameNode restart timeout parameter, `upgrade.parameter.nn-restart.timeout` in `/etc/ambari-server/conf/ambari.properties` on the Ambari Server host. You may need to add the restart timeout parameter and value to the Ambari server host following a default installation. For a large cluster, you should add 10% to the usual time (in seconds) required to restart your NameNode. Although no standard way to determine an appropriate value exists, you may use the following guidance:

For example, record the time (seconds) required to restart the active NameNode for your current Ambari server version. If restarting takes 10 minutes, (600 seconds), then add `upgrade.parameter.nn-restart.timeout=660` to the `/etc/ambari-server/conf/ambari.properties` file on the Ambari Server host.

After adding or resetting the Ambari NameNode restart parameter, restart your Ambari server before starting the HDP upgrade.

`ambari-server restart`

For Ambari Upgrades

- You must review the [Behavioral Changes](#).
- Review supported Ambari Server database versions using the *Cloudera Support Matrix*. Review the resources required for the database upgrade. You will upgrade your Ambari Server database during the Ambari upgrade procedure.
- Ambari 7.1.x.x only supports the following operations when running against a HDP 2.6.5.x cluster:
  - Run Service Checks
  - Start, Stop, Restart a Service
  - Change Configuration
  - Enable and Disable Maintenance Mode
  - Disable Auto Start
  - Remove Services
  - Remove Components
  - Remove and Decommission Hosts

To restore full management functionality, use Ambari 7.1.x.x to upgrade to HDP intermediate bits as soon as possible.

For HDP Cluster Upgrades

- Ensure sufficient disk space on /usr/hdp/<version> (roughly 10GB for each additional HDP release).
- Additional components are added to the cluster as part of the HDP intermediate bits upgrade including YARN ATSV2, YARN Registry DNS, and additional Hive clients required for the Spark history server. If your cluster has kerberos enabled, you must configure Ambari to manage the kerberos administrator credentials prior to the upgrade, so that the appropriate kerberos principals can be created during the upgrade process.
- You must take a backup of the running topology processes if your HDP cluster includes the Storm component. You must stop the Storm services during the upgrade. Cloudera Private Cloud Base cluster does not support Storm.

### Related Information

[Cloudera Support Matrix](#)

## Kerberos cluster

To successfully install the HDP intermediate bits using Ambari as a part of the CDP upgrade process, you must review and configure your KDC administrator credentials, administrator host FQDN, and kadmin principal.



**Important:** For a Kerberized cluster, you must review and configure your KDC:

- KDC Admin Credentials:

The Ambari Server adds new components as a part of the HDP 2.6.5.x to HDP intermediate bits upgrade and needs to be configured to enable password encryption and saving the KDC admin credentials so that necessary principals can be created. For steps on how to do this, see [Encrypt Database and LDAP Passwords in Ambari](#), [Managing Admin Credentials](#), and [Authentication and Kerberos Issues](#).

- KDC Admin Host FQDN:

Ensure that the KDC Admin host is set to a fully qualified domain name (FQDN) and not an IP address. You can check this setting by going to Services > Kerberos > Configs > KDC Hosts

- KDC kadmin principal:

The MIT KDC integration in Ambari 7.1.x.x has been improved to connect more securely with the Kerberos Administration server (kadmind). This increased security expects a Kerberos admin service principal to be present with a specifically formatted principal name. The format expected is kadmin/fully.qualified.kdc.hostname@REALM. This expectation is a behavior change from previous versions of Ambari, where having such a Kerberos admin service principal was not required. This principal is present by default in most MIT KDC installations, but some customers have reported that this principal does not exist in their KDC. Due to this, it is recommended that before upgrading a kerberized cluster to Ambari 7.1.x.x, you ensure this principal exists in your KDC database.

## Before upgrading any cluster

Prior to installing HDP intermediate bits using Ambari as a part of the CDP upgrade process, your cluster must meet the certain prerequisites to ensure that the cluster is in a healthy operating mode and can successfully manage the upgrade process.

For any Cluster

**Disk Space:** Be sure to have adequate space on /usr/hdp for the target HDP version. Each complete install of an HDP version occupies about 5 GB of disk space.

**Ambari Agent Heartbeats:** All Ambari Agents must be communicating and heartbeating to Ambari Server. Any hosts that are not heartbeating must be in Maintenance mode.

Hive Upgrade: The upgrade process does not back up the Hive Metastore, nor does it compact ACID tables. Before upgrading Hive, you must:

- Manually back up your Hive metastore database after using the pre-upgrade tool, described later, and before upgrading.
- If you have ACID tables in your Hive metastore database, enable ACID operations using Ambari Web or set Hive configuration properties to enable ACID.



**Note:** You must remove the components that are unsupported in CDP Private Cloud Base. For more information, see [HDP Core component version changes](#). If you are not using a component in the CDP cluster, then you must remove them. The unused component extends the time and complexity of the HDP to CDP upgrade.

Host Maintenance Mode

The following two scenarios are checked:

- Any hosts in Maintenance mode must not be hosting any Service Master Components.
- Hosts running Worker Components display a warning if they are in Maintenance mode. You can proceed with the upgrade. But these hosts will not be upgraded. Before finalising the upgrade, you must delete these hosts from the cluster.

Service Maintenance Mode: No services can be in Maintenance mode, except for Ambari Metrics System, SmartSense, and Log Search.

Services Started: All services must be started, except for Ambari Metrics System, SmartSense, and Log Search.

Service Checks: All service checks must pass. Be sure to run Service Actions > Run Service Check on all services (and remediate if necessary) prior to attempting an HDP upgrade.

All Prior Upgrades are Finalized: Any prior upgrade started with Ambari must be finalized.

## Managing MPacks

You must review the information mentioned in this section before you proceed with the upgrade.

- Not all Management Packs are compatible with the HDP intermediate bits. Cloudera Support can help you assess the feasibility of your upgrade if you have Management Packs other than HDF installed.
- Cloudera does not support HDF in-place upgrades if HDF is used in combination with the HDP 2 cluster. This is because it is not supported to have HDP 2.x + HDF 3.5.2 managed by the same Ambari.
- If you're running HDF on top of HDF 2.6.5, you cannot upgrade HDF. You have two options:
  - You can upgrade both HDP and HDF in order to have Ambari 2.7.5 managing both HDP 3.1.5 and HDF 3.5.2. You can then proceed with the in-place upgrade to CDP.
  - You can detach the HDF services from Ambari, uninstall the HDF mpack from Ambari, and proceed with the upgrade for Ambari/HDP. After upgrading to CDP, you can install the corresponding services and migrate the workloads from the HDF cluster to the new services running on CDP.



**Note:** NiFi and NiFi Registry are part of the CFM parcel needs to be installed in CDP. Kafka and Schema Registry are part of the CDP runtime. Storm and SAM are deprecated and replaced by Flink and SSB. For more information, see [CFM Migration](#) on migrating NiFi flows between two environments.

- If you have any doubts or concerns or questions about the migration process, you must contact the Cloudera Support for further assistance to help you determine the best migration plan for your business needs.
- The latest and supported version of AM2CM is AM2CM 2.0.2. This version is capable of handling the HDF 3.5.2.0 # CPD 7.1.7.0 upgrade path.

## 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 . However, you can add Accumulo as a service on the managed CDP Private Cloud Base.

**Important:**

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

## 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          | (CM)          |
|------------|-----------------------|-------------------------|---------------|
| Hadoop     | 2.7.3                 | 3.1.1                   | 3.1.1         |
| Hive       | 1.2.1<br>2.1.0 (LLAP) | 3.1.3000<br>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<br>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.

## Upgrading the cluster's underlying OS

Ensure that all your hosts in the cluster are on the operating systems supported with the HDP intermediate bits and Ambari 7.1.x.x before starting the upgrade from HDP 2.6.5.x to HDP intermediate bits.



Only RHEL, CentOS, and Ubuntu operating systems are supported with the HDP intermediate bits and Ambari 7.1.x.x. Ensure that all your hosts in the cluster are on the supported operating system before starting the upgrade from HDP 2.6.5.x to HDP intermediate bits. For more information on the supported versions of Operating systems, see [Operating system requirements](#).



**Note:** SLES 12 SP5 is now supported for use with the HDP intermediate bits and Cloudera Private Cloud Base 7.1.4 and higher.

For many, this is a process that takes time and orchestration between multiple teams within your organization. Two high-level guidelines for moving from one major operating system version to another is as follows:

In-Place and Restore:

Perform an In-place OS refresh and use Ambari Restore Host feature

Move and Decom:

Move Masters and Decom/Recom Workers

Each option has advantages and disadvantages and high-level decision criteria.

## In-Place and Restore

Perform an In-place OS refresh and use Ambari Restore Host feature.

This option should be used in medium to large clusters (25 or more nodes), with operational teams that have environment automation experience, and have followed best practices when setting up component High Availability and HDP directory structures (such as ensuring that the HDP component data and metadata are not stored on the root volume).

This option involves going through each host in the cluster and ensuring important metadata and data are stored on a volume that is not being used by the operating system, and leverages component high availability to maintain maximum cluster availability. When visiting each host, the host is shut down, the operating system volume is refreshed with the new version of the chosen operating system, the host is configured with the same IP address and hostname, all volumes are re-mounted, and the Ambari agent is installed and configured. After the host has rejoined the cluster, the Ambari Recover Host functionality is used to reinstall, reconfigure, and start services. To ensure that no data is lost during the reinstall of the operating system, verify that your OS volumes do not contain any HDP data or metadata. Additionally, during the OS reinstall, make sure that you do not erase or reformat any non-operating-system volumes, such as HDFS data drives, as this may result in data loss.

## Move and Decommission

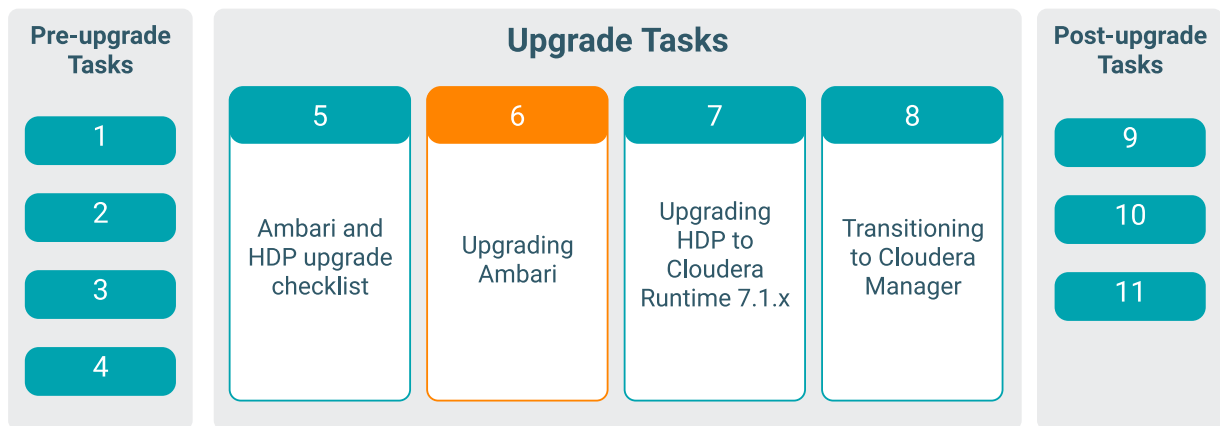
Move Masters and Decom/Recom Workers. Each option has advantages and disadvantages and high-level decision criteria.

This option should be used in smaller clusters (under 25 nodes), where operational teams may not have access to operating system and configuration management automation tooling or have not yet followed best practices when setting up the HDP directory structures (such as ensuring HDP component data and metadata are not stored on the root volume).

This option involves decommissioning worker nodes and replacing them with worker nodes that have the new operating system version on them. For master nodes, the move-master operation is used to move all masters off of a host, and on to a new host with the new operating system version on them. Decommissioning worker nodes can take a great deal of time, depending on the density of the nodes, and move-master operations require many cluster services to be restarted, so this is a time-consuming process that requires multiple periods downtime, but it does not require any operating system level operations to be performed.

# Upgrading Ambari

You must understand the considerations and the process to upgrade from the current version of Ambari to Ambari 7.1.x.x.



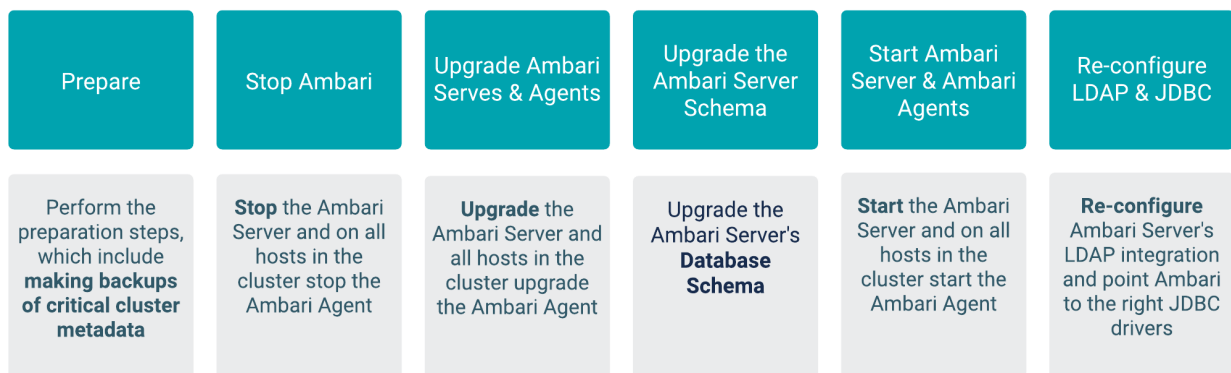
## Before you upgrade Ambari

You must be aware of the exact steps and understand the interdependencies and order of the steps so that you can adjust and account for any special configurations for your cluster before starting the Ambari upgrade process.

Upgrading Ambari is much like any other upgrade in HDP. The main difference here is that the underlying platform is running Cloudera Runtime bits for all services. This is important because it provides an upgrade path to using the latest bits of the underlying stack. These include upgrades to HDFS, Hive, Spark, Ranger, Oozie, HBase, and so on.

Ambari and the HDP intermediate bits clusters being managed by Ambari can be upgraded independently. This section describes the process to upgrade Ambari. You are encouraged to read completely through this entire document before starting the upgrade process, so that you understand the interdependencies and order of the steps. It is recommended you validate these steps in a test environment to adjust and account for any special configurations for your cluster.

- Preparing to Upgrade Ambari
- Upgrade Ambari
- The high-level process for upgrading Ambari is as follows:



**Note:** Ambari 7.1.x.x is not intended to be a long-term solution, only a brief interim step to CDP Private Cloud Base.



**Important:** If you stay on HDP 7.1.x you will not be able to upgrade directly to HDP 7.1.y (where  $y > x$ ). You must always first upgrade from HDP 7.1.x to CDP 7.1.x and then upgrade to CDP-7.1.y.

## Setting up a local repository

You need to set up a local repository, update the version repository base URLs, and edit the repository configuration file.

As a first step, you must set up a local repository for Ambari and HDP. For more information, see [set up a local repository](#).

A case-study of setting up local repositories: [Case study for setting up a local repo](#)

When the local repository is created, update the version repository base urls in Ambari. For more information, see [update the version repository base urls](#).

Finally, edit the repository configuration file to use this new local repository. For more information, see [edit the repository configuration file](#)



**Note:** You can use the case studies to set up a local repository for all of the 7.1.x versions. For example, the case studies in this section help you to set up a local repository for the 7.1.7 version.

## Updating Ambari repo files

Some services depend on components that are installed from the Ambari repository. It is not updated automatically. Also, Package Manager displays an error about the unavailable repository URL when you update the package lists. Cloudera recommends you to manually update the URLs located at `/etc/yum.repos.d/ambari.repo` on all hosts (including server host).

## Updating HDP repo files

When the cluster settings for the HDP repository URL is updated, repository files on hosts are not immediately regenerated. The files are re-generated when you add a new component or service. But an inaccessible repository URL causes the Package Manager to display an error about the unavailable repository URL when you update the package lists. Cloudera recommends you to manually update the URLs in the HDP repository files (for example, `/etc/yum.repos.d/ambari-hdp-1.repo`) on all agent hosts.

### Before you begin



**Note:** Ambari 7.1.x.x and HDP 7.1.x.x are version numbers. For example, Ambari 7.1.8.0 and HDP 7.1.8.0 or Ambari 7.1.7.0 and HDP 7.1.7.0 or Ambari 7.1.6.0 and HDP 7.1.6.0 and so on.

## Case study for setting up an HDP-GPL local repository

If the cluster is using the GPL components, you must replace the HDP-GPL repository URL with a local repository.

### Procedure

1. Set up a local HDP-GPL repository

```
wget -nv https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/centos7/hdp.repo -O /etc/yum.repos.d/hdp.gpl.repomkdir -p /var/www/html/hdp-gpl/centos7cd /var/www/html/hdp-gpl/centos7/reposync -r HDP-GPL-7.1.x.xcreaterepo /var/www/html/hdp-gpl/centos7/HDP-GPL-7.1.x.x/
```

2. Edit `/etc/ambari-server/conf/ambari.properties`.
3. Replace `gpl.license.accepted=false` with `gpl.license.accepted=true`.
4. Restart Ambari server.
5. You must edit the HDP-GPL repository similar to the HDP repository. However, the URL and Local repository contents, and the UI fields are different for both the repositories. For more information on updating the version repository, see [Update version repository base urls](#). If you do not plan use the GPL components, disable the `gpl.license.accepted` property.

## 6. Restart Ambari server

### Setting up local repository with temporary internet access


You must have completed the Getting Started Setting up a Local Repository procedure.

#### About this task

To finish setting up your local repository, complete the following:

#### Procedure

1. Install the repository configuration files for Ambari and the Stack on the host.
2. Confirm repository availability:
  - For RHEL, CentOS, Oracle or Amazon Linux: `yum repolist`
  - For SLES: `zypper repos`
  - For Debian and Ubuntu: `dpkg-list`
3. Synchronize the repository contents to your mirror server:
  - a) Browse to the web server directory:
    - For RHEL, CentOS, Oracle or Amazon Linux: `cd /var/www/html`
    - For SLES: `cd /srv/www/htdocs/rpms`
    - For Debian and Ubuntu: `cd /var/www/html`
  - b) For Ambari, create the ambari directory and `reposync`: `mkdir -p ambari/<OS>`, `cd ambari/<OS>`, and `reposync -r Updates-Ambari-7.1.x.x`. In this syntax, the value of `<OS>` is `amazonlinux2`, `centos7`, `sles12`, `ubuntu16`, `ubuntu18`, or `debian9`.
 

 **Important:** Due to a known issue in version 1.1.31-2 of the Debian 9 `reposync`, we advise using `reposync` version 11.3.1-3 or above when working on a Debian 9 host.
  - c) For Hortonworks Data Platform (HDP) stack repositories, create the `hdp` directory and `reposync`: `mkdir -p hdp/<OS>`, `cd hdp/<OS>`, and `reposync -r HDP-<latest.version>`, and `reposync -r HDP-UTILS-<version>`.
  - d) For HDF Stack Repositories, create an `hdf` directory and `reposync`. `mkdir -p hdf/<OS>`, `cd hdf/<OS>`, and `reposync -r HDF-<latest.version>`.
4. Generate the repository metadata:
  - For Ambari: `createrepo <web.server.directory>/ambari/<OS>/Updates-Ambari-7.1.x.x`
  - For HDP Stack Repositories: `createrepo <web.server.directory>/hdp/<OS>/HDP-<latest.version>` `createrepo <web.server.directory>/hdp/<OS>/HDP-UTILS-<version>`
  - For HDF Stack Repositories: `createrepo <web.server.directory>/hdf/<OS>/HDF-<latest.version>`

5. Confirm that you can browse to the newly created repository:

- Ambari Base URL `http://<web.server>/ambari/<OS>/Updates-Ambari-7.1.x.x`
- HDF Base URL `http://<web.server>/hdf/<OS>/HDF-<latest.version>`
- HDP Base URL `http://<web.server>/hdp/<OS>/HDP-<latest.version>`
- HDP-UTILS Base URL `http://<web.server>/hdp/<OS>/HDP-UTILS-<version>`

Where:

- `<web.server>` – The FQDN of the web server host
- `<version>` – The Hortonworks stack version number
- `<OS>` – `centos7`, `sles12`, `ubuntu16`, `ubuntu 18`, or `debian9`



**Important:** Be sure to record these Base URLs. You need them when you are installing Ambari and the Cluster.

6. Optional. If you have multiple repositories configured in your environment, deploy the following plug-in on all the nodes in your cluster.

- Install the plug-in. For RHEL/CentOS/Oracle 7: `yum install yum-plugin-priorities`
- Edit the `/etc/yum/pluginconf.d/priorities.conf` file to add the following: `[main] enabled=1 gpgcheck=0`

## Case study for setting up local repository

Review this case study to understand - how to prepare a local repository for Ambari 7.1.x.x and HDP 7.1.x.x on Centos7:

```

yum install yum-utils createrepo -y
yum install httpd -y
//firewall configuration
sudo systemctl start httpd
sudo systemctl status httpd
mkdir -p /var/www/html/
wget -nv https://archive.cloudera.com/p/ambaridc/7.x/7.1.7.2000/centos7/amb
aridc.repo -O /etc/yum.repos.d/ambari.repo
wget -nv archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/centos7/hdp.repo -O /
etc/yum.repos.d/hdp.repo
yum repolist
cd /var/www/html
mkdir -p ambari/centos7
cd ambari/centos7/
reposync -r ambari-7.1.x.x
cd ../../
mkdir -p hdp/centos7
cd hdp/centos7/
reposync -r HDP-7.1.x.x
reposync -r HDP-UTILS-1.1.0.22
createrepo /var/www/html/ambari/centos7/ambari-7.1.x.x
createrepo /var/www/html/hdp/centos7/HDP-7.1.x.x
createrepo /var/www/html/hdp/centos7/HDP-UTILS-1.1.0.22/

```

Result:

The repositories will be available at the local web server:  
`http://<web.server>/ambari/centos7/ambari-7.1.x.x/`  
`http://<web.server>/hdp/centos7/HDP-7.1.x.x/`  
`http://<web.server>/hdp/centos7/HDP-UTILS-1.1.0.22/`

## Update version repository base urls

Use Admin/Versions/Repositories to provide URLs for your source repositories.

### Procedure

1. Browse to Ambari Admin > Versions > Manage Versions > OK.
2. In Admin/Versions, click the version name you want to modify.
3. In Repositories, modify the base URLs for the repositories. To use the private software repositories, see the list of available HDP repositories for each OS. Or, if you are using a local repository, enter the Base URLs for the local repository you have created.
4. Click Save.
5. Click Confirm Change. You must confirm the change since you are about to change repository Base URLs that are already in use. Please confirm that you intend to make this change and that the new Base URLs point to the same exact stack version and build.

## Preparing Ambari Repository Configuration File to use Local Repository

You must prepare the Ambari Repository Configuration File to use the Local Repository.

### Procedure

1. Download the ambari.repo file from the private repository: <https://username:password@archive.cloudera.com/p/ambaridc/7.x/7.1.7.2000/<OS>/ambaridc.repo/>
2. Edit the ambari.repo file and replace the Ambari Base URL baseurl obtained when setting up your local repository.

```
#VERSION_NUMBER=7.1.7.2000-3
[ambari-7.1.7.2000-3]
#json.url = http://public-repo-1.hortonworks.com/HDP/hdp_urlinfo.json
name=ambari Version - ambari-7.1.7.2000
baseurl=https://archive.cloudera.com/p/ambaridc/7.x/7.1.7.2000/centos7/
gpgcheck=1
gpgkey=https://archive.cloudera.com/p/ambaridc/7.x/7.1.7.2000/centos7/RPM-GPG-K
G-K
KEY/RPM-GPG-K
EY-Jenkins
enabled=1
priority=1
```



**Note:** You can disable the GPG check by setting `gpgcheck=0`. Alternatively, you can keep the check enabled but replace `gpgkey` with the URL to GPG-KEY in your local repository.

3. Base URL for a Local Repository.
  - Built with Repository Tarball (No Internet Access) `http://<web.server>/Ambari-7.1.x.x/<OS>`
  - Built with Repository File (Temporary Internet Access) `http://<web.server>/ambari/<OS>/Updates-Ambari-7.1.x.x` where `<web.server>` = FQDN of the web server host, and `<OS>` is `amazonlinux2`, `centos7`, `sles12`, `ubuntu16`, `ubuntu18`, or `debian9`.
4. Place the ambari.repo file on the host you plan to use for the Ambari server:
  - For RHEL/CentOS/Oracle/Amazon Linux: `/etc/yum.repos.d/ambari.repo`
  - For SLES: `/etc/zypp/repos.d/ambari.repo`
  - For Debian/Ubuntu: `/etc/apt/sources.list.d/ambari.list`
5. Edit the `/etc/yum/pluginconf.d/priorities.conf` file to add the following values: `[main] enabled=1 gpgcheck=0`

## Backup Ambari

You must back up the current Ambari database. It creates a copy of the current Ambari database and meta information.

## Procedure

1. On the Ambari Server host, stop the original Ambari Server.

```
ambari-server stop
```

2. Create a directory to hold the database backups.

```
cd /tmp
mkdir dbdumps/
cd dbdumps/
```

3. Create the database backups.

### POSTGRES

```
pg_dump [AMBARI_DB_USERNAME] -U [AMBARI_DB_USERNAME] -f pre_upgrade.sql
Password: [AMBARI_DB_PASSWORD]
```

### MYSQL

```
mysqldump [AMBARI_DB_NAME] -u [AMBARI_DB_USERNAME] > pre_upgrade.sql
```

| Variable             | Description            | Default |
|----------------------|------------------------|---------|
| [AMBARI_DB_NAME]     | The database name.     | ambari  |
| [AMBARI_DB_USERNAME] | The database username. | ambari  |
| [AMBARI_DB_PASSWORD] | The database password. | bigdata |

4. Create a backup of the Ambari Server properties.

```
/etc/ambari-server/conf/ambari.properties
/etc/ambari-agent/conf/ambari-agent.ini
```

## Ambari Behavioral changes

Behavioral changes denote a marked change in behavior from the previously released version to this version of Ambari. Between Ambari 2.6.x and Ambari 7.1.x there were changes implemented in 2.7.x versions.

### Ambari 7.1.x

| Description  | Behavioral Changes  |
|--|---|
| Atlas and Ranger Default Installation                          | During the service selection process of a new install, you see multiple limited functionality warnings if Atlas, Ranger, and their dependent services are not selected for installation.  |
| Ambari UI uses websocket to communicate with the Ambari Server | The Ambari UI now uses websocket to communicate with the Ambari Server. If you are using Ambari through a proxy like Nginx, Knox, or Apache HTTPD, ensure that that proxy is configured to proxy websocket traffic. If this is not done, the Ambari web UI does not update quickly and performance is sluggish. |
| Ambari Server and Log Search UI Changes                        | The look and feel of the Ambari Server UI in Ambari 7.1.x.x and the Log Search UI have changed significantly.   |
| Ambari Server Setup SSO CLI                                    | The ambari-server setup-sso command now includes a prompt to setup SSO for Ambari, Atlas, and Ranger.   |

|   |   |
|---|---|
| KDC Admin host must use FQDN                      | As of Ambari 7.1.x, the KDC admin host value is expected to be the fully qualified domain name (FQDN) of the host where the KDC administration server is located. If Kerberos is enabled, then verify and update the KDC Admin host to use FQDN. This must be done before the upgrade can proceed.  |
| Keytab Regeneration                               | After upgrading to Ambari 7.1.x, clicking Regenerate Keytabs and selecting only to regenerate missing keytabs regenerates all keytabs. This behavior only happens the first time this operation is performed, all subsequent use of Regenerate Keytabs with only missing keytabs selected behaves as usual.   |
| KDC Admin Credentials Needed during Stack upgrade | When upgrading from HDP 2.6.5.x to HDP intermediate bits, additional YARN components are added to the cluster. If your cluster has Kerberos enabled, the Ambari Server must be configured to store the KDC admin credentials so that the principals for these components can be created during the upgrade.   |
| Blueprint Deletion                                | You can not allow a DELETE operation on an existing Blueprint when a cluster deployment has completed with that Blueprint. In any other case, the DELETE is allowed.  |
| Version API Change                                | The default response to the /api/v1/clusters has changed. Before Ambari 2.7, each cluster's "version" was reported in the response of this call by default. In Ambari 7.1.x, this "version" field is no longer included in the default response, but can be included in the response if the following query parameters are used: "/api/v1/clusters?fields=Clusters/version" |
| Ambari Views Removed                              | Ambari 7.1.x does not support Hive, Tez, and Pig Views.   |

| Cloudera Bug ID | Apache component | Apache JIRA | Summary  | Details   |
|-----------------|------------------|-------------|--|---|
| BUG-105818      | Ambari           | AMBARI-9016 | HTTP Header Content-Type:text/plain when the content is in JSON format | Scenario: Ambari REST API is expecting "Content-Type" HTTP header set to "text/plain" even though content is in JSON format. If the client specified the same header with the value of "application/json", Ambari threw an HTTP 500 error. Previous Behavior: Ambari behaviour was misleading, forcing the user to use "text/plain" Content-Type header in HTTP clients. New Behavior: Ambari expects an "application/json" Content-Type header. Expected Customer Action: Revise any custom code that uses the Ambari REST API and take actions if needed. |

## Ambari Properties backup

Prior to upgrading from Ambari 2.6.2.x to Ambari 7.1.x.x, you must make a copy of the Ambari Server configuration file so that you can reapply any configuration changes you have done earlier to the new file.

During the Ambari upgrade, the existing /var/lib/ambari-server/ambari-env.sh file is overwritten and a backup copy of ambari-env.sh (with extension .rpmsave) is created. If you have manually modified ambari-env.sh (for example, to change Ambari Server heap), you must reapply your changes in the new file.

To purge your Ambari cluster, see [Purging Ambari Server Database History](#).



## Review Ambari UI and the Quick Links

Ensure to review the Ambari UIs and quick links prior to the upgrade. This check is helpful to compare the user interface state post upgrade.

You must:

- Review each page from Ambari and take a screenshot for future reference.
- Review each quick link in Ambari and take a screenshot for future reference. For example, Solr UI, Ranger UI, and RM UI, etc.
- Check all services are functional on Ambari side.

## Upgrade to Ambari 7.1.x.0

After reviewing the information related to Ambari UI and the Quick Links, and backing up the Ambari Server configuration file, perform the recommended steps for upgrading Ambari.

### Upgrade to Ambari

## Upgrade to Ambari 7.1.7 SP2 (7.1.7.2000-3)

After reviewing the information related to Ambari UI and the Quick Links, and backing up the Ambari Server configuration file, perform the recommended steps for upgrading Ambari.

### Procedure

1. Turn off Auto Restart from Ambari web UI by browsing to Admin > Service Auto Start. Set Auto-Start Services to Disabled. Click Save.
2. If you are running SmartSense in your cluster, stop the service. From Ambari Web, browse to Services > SmartSense and select Stop from the Service Actions menu. Then, select Turn on Maintenance Mode from the Service Actions menu.

3. If you are running Ambari Metrics in your cluster, put it in Maintenance Mode. Don't stop the service yet. From Ambari Web, browse to Services > Ambari Metrics and select Turn on Maintenance Mode from the Service Actions menu.

a) Download the pre-upgrade script

1. SSH into the host running the Metrics Collector. Note, that in case Ambari Metrics System (AMS) is running in High Availability mode, there are multiple hosts running the Metrics Collector. These steps must be done on each Metrics Collector host.
2. Become the AMS service user. For example, run the following su command on Linux: `$ sudo su - ams`
3. Change to the /tmp directory.
4. Execute the following command to download the pre-upgrade script for CentOS 7:

```
$ wget https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/7.x/7.1.7.2000/centos7/tars/ambari/ambari-metrics-pre-upgrade-7.1.7.2000.3.tar.gz
```

Execute the following command to download the pre-upgrade script for Ubuntu 18:

```
$ wget https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/7.x/7.1.7.2000/ubuntu18/tars/ambari/ambari-metrics-pre-upgrade-7.1.7.2000.3.tar.gz
```

Execute the following command to download the pre-upgrade script for SLES 12:

```
$ wget https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/7.x/7.1.7.2000/sles12/tars/ambari/ambari-metrics-pre-upgrade-7.1.7.2000.3.tar.gz
```

To find all the download links in a centralised location, see [Software download matrix](#)

b) Extract and run the pre-upgrade script

1. Extract the pre-upgrade tarball: `$ tar xvzf ambari-metrics-pre-upgrade-<version>.tar.gz`
2. Change to the ambari-metrics-pre-upgrade-<version> directory. `$ cd ambari-metrics-pre-upgrade-<version>`
3. Make the pre-upgrade script executable: `$ chmod +x ams-pre-upgrade-2.6-to-7.x-<version>.sh`
4. Run the pre-upgrade script on all Metrics Collector hosts: `$ ./ams-pre-upgrade-2.6-to-7.x-<version>.sh`

c) Stop the service. From Ambari Web, browse to Services > Ambari Metrics and select Stop from the Service Actions menu.

4. If you are running Log Search in your cluster, stop the service. From Ambari Web, browse to Services > Log Search and select Stop from the Service Actions menu. Then, select Turn on Maintenance Mode from the Service Actions menu.
5. Stop the Ambari Server. On the host running Ambari Server: `ambari-server stop`
6. Stop all Ambari Agents. On each host in your cluster running an Ambari Agent: `ambari-agent stop`.
7. Back up your existing Ambari Server database. For example, to back up the default, embedded postgres db for Ambari:

```
mkdir /root/backups/
pg_dump -U ambari ambari > /root/backups/ambari-before-upgrade.sql
```

The default password is: bigdata.



**Note:** You must take a back up of the Ambari database if you have hosted the Ambari database on any other supported databases like MySQL or MariaDB.

8. Upgrade your Ambari Server database to a current supported version, only if your Ambari Server database version is not supported on the target Ambari version. To verify current supported versions, see the *Cloudera Support*

*Matrix*. For example, if you are running postgres v9.4, you must upgrade postgres to v9.6 or v10.2. If you are running MariaDB 10.2, upgrading the Ambari Server database is not required.

See the documentation for your database type and version for specific database upgrade instructions.

- Retrieve the new Ambari repo and replace the old repository file with the new repository file on all hosts in your cluster.



**Important:** Check your current directory before you download the new repository file to make sure that there are no previous versions of the `ambaridc.repo` file. If you do not, and a previous version exists, the new download is saved with a numeric extension, such as `ambari.repo.1`. Make sure that you copy the new version.

Select the repository appropriate for your environment from the following list:

- For RHEL or CentOS:

```
wget -nv https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/7.x/7.1.7.2000/centos7/ambari.repo -O /etc/yum.repos.d/ambari.repo
```

- For Ubuntu:

```
wget -nv https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/7.x/7.1.7.2000/ubuntu18/ambari.list -O /etc/apt/sources.list.d/ambari.list
```

- For Sles:

```
wget -nv https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/7.x/7.1.7.2000/sles12/ambari.repo -O /etc/zypp/repos.d/ambari.repo
```

- Upgrade Ambari Server. On the host running Ambari Server:

For RHEL or CentOS:

```
yum clean all
```

```
yum info ambari-server
```

In the information output, visually validate that there is an available version containing "7.1.7.2000"

```
yum upgrade ambari-server
```

For Ubuntu 18:

```
apt-get clean all
```

Before upgrading Ambari Server, you must update the username and password in the `ambari.list` file. Run the following command:

```
vi /etc/apt/sources.list.d/ambari.list
```

For example, the output displays the following:

```
#VERSION_NUMBER=7.1.7.2000-3
#json.url = https://archive.cloudera.com/p/HDP/hdp_urlinfo.json
deb https://archive.cloudera.com/p/ambaridc/7.x/7.1.7.2000/ubuntu18 Ambari main
apt-get update
```

```
apt-cache show ambari-server | grep Version
```

In the information output, visually validate that there is an available version containing "7.1.7.2000"

```
apt-get upgrade ambari-server
```

For SLES 12:

```
zypper clean
```

Before upgrading the Ambari Server, you must create, place, and set the credentials for zypper. Run the following command:

```
vi /etc/zypp/repos.d/ambari.repo
```

Add the following postfix to the end of BaseURL and gpgkey: ?credentials=ambari.cat For example, the output displays the following:

```
#VERSION_NUMBER=7.1.7.2000-3
[ambari-7.1.7.2000]
#json.url = https://archive.cloudera.com/p/HDP/hdp_urlinfo.json
name=ambari Version - ambari-7.1.7.2000
baseurl=https://archive.cloudera.com/p/ambaridc/7.x/7.1.7.2000/sles12
gpgcheck=0
gpgkey=https://archive.cloudera.com/p/ambaridc/7.x/7.1.7.2000/sles12/RPM-GPG-KEY/RPM-GPG-KEY-Jenkins
enabled=1
priority=1
```

Create the following file with your paywall credentials (replace <username> with your paywall username and <password> accordingly): /etc/zypp/credentials.d/ambari.cat

```
username=<username>
password=<password>
```

```
zypper info ambari-server
```

In the information output, visually validate that there is an available version containing "7.1.7.2000"

```
Zypper up ambari-server
```

**11.** Check for upgrade success by noting progress during the Ambari Server installation process you started in Step 10.

- As the process runs, the console displays output similar to the following:  
Setting up Upgrade Process Resolving Dependencies Running transaction check
- If the upgrade fails, the console displays output similar to the following:  
Setting up Upgrade Process No Packages marked for Update
- A successful upgrade displays output similar to the following:  
Updated: ambari-server.noarch 0:7.1.7.2000. Complete!

**12.** Upgrade all Ambari Agents. On each host in your cluster running an Ambari Agent:

- For RHEL or CentOS: Before upgrading Ambari Agent, you must update the username and password in the `ambaridc.repo` file. Run the following command:

```
vi /etc/yum.repos.d/ambaridc.repo
```

For example, the output displays the following:

```
#VERSION_NUMBER=7.1.7.2000-3
[ambari-7.1.7.2000]
#json.url = https://archive.cloudera.com/p/HDP/hdp_urlinfo.json
name=ambari Version - ambari-7.1.7.2000
baseurl=https://archive.cloudera.com/p/ambaridc/7.x/7.1.7.2000/centos7
gpgcheck=1
gpgkey=https://archive.cloudera.com/p/ambaridc/7.x/7.1.7.2000/centos7/
RPM-GPG-KEY/RPM-GPG-KEY-Jenkins
enabled=1
priority=1
```

```
yum upgrade ambari-agent
```

- For Ubuntu:

```
apt-get update
```

```
apt-get upgrade ambari-agent
```

- For SLES: Copy the following files from the server host to each host with the same path:

```
/etc/zypp/credentials.d/ambari.cat
/etc/zypp/repos.d/ambari.repo
```

Then run the following command:

```
zypper up ambari-agent
```

**13.** After the upgrade process completes, check each host to make sure the new files have been installed:

For CentOS 7.7:

```
rpm -qa | grep ambari-agent
```

For Ubuntu 18:

```
dpkg -l ambari-agent
```

**14.** Upgrade Ambari Server database schema. On the host running Ambari Server: `ambari-server upgrade`. When the Ambari Server database schema has been upgraded, you should see command output like this: Ambari Server ‘upgrade’ completed successfully.**15.** Start the Ambari Server. On the host running Ambari Server: `ambari-server start`**16.** Start all Ambari Agents. On each host in your cluster running an Ambari Agent: `ambari-agent start`**17.** Open Ambari Web UI. Point your browser to the Ambari Web UI:

- When Ambari Server is configured for HTTPS: `https://<your.ambari.server>:8443`
- When Ambari Server is configured for HTTP: `http://<your.ambari.server>:8080`

where `<your.ambari.server>` is the name of your ambari server host. For example, `c7401.ambari.apache.org`.



**Important:** Refresh your browser so that it loads the new version of the Ambari Web code. If you have problems, clear your browser cache manually, then restart Ambari Server.

18. Log in using the Ambari administrator credentials that you have set up. For example, the default name/password is admin/admin. You will see a Restart indicator next to each service after upgrading. Ambari upgrade has modified the configuration properties of your cluster based on the new configuration types and properties being made available for each service with this release of Ambari. Review these changes by comparing the previous configuration with the latest version created by "ambari-upgrade".



**Caution:** Do not manually restart these services unless future steps in the upgrade guide prompt you to do so. Manually restarting these services may significantly disrupt your upgrade. Ambari will restart each service automatically during the HDP upgrade.

19. If you have configured Ambari to authenticate against an external LDAP or Active Directory, you must re-run `ambari-server setup-ldap`.
20. If you have configured your cluster for Hive, Ranger or Oozie with an external database (Oracle, MySQL or PostgreSQL), you must run the following command to get the JDBC driver jar file in place: `ambari-server setup --jdbc-db` and `--jdbc-driver`.  
A sample command to setup MySQL connector: `ambari-server setup --jdbc-db=mysql --jdbc-driver=/path/to/mysql/mysql-connector-java.jar`

### Upgrade to Ambari 7.1.7.1

After reviewing the information related to Ambari UI and the Quick Links, and backing up the Ambari Server configuration file, perform the recommended steps for upgrading Ambari.

#### Procedure

1. Turn off Auto Restart from Ambari web UI by browsing to Admin > Service Auto Start. Set Auto-Start Services to Disabled. Click Save.
2. If you are running SmartSense in your cluster, stop the service. From Ambari Web, browse to Services > SmartSense and select Stop from the Service Actions menu. Then, select Turn on Maintenance Mode from the Service Actions menu.

3. If you are running Ambari Metrics in your cluster, put it in Maintenance Mode. Don't stop the service yet. From Ambari Web, browse to Services > Ambari Metrics and select Turn on Maintenance Mode from the Service Actions menu.

a) Download the pre-upgrade script

1. SSH into the host running the Metrics Collector. Note, that in case Ambari Metrics System (AMS) is running in High Availability mode, there are multiple hosts running the Metrics Collector. These steps must be done on each Metrics Collector host.
2. Become the AMS service user. For example, run the following su command on Linux: `$ sudo su - ams`
3. Change to the /tmp directory.
4. Execute the following command to download the pre-upgrade script for CentOS 7:

```
$ wget https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/7.x/7.1.7.1/centos7/tars/ambari/ambari-metrics-pre-upgrade-7.1.7.1.14.tar.gz
```

Execute the following command to download the pre-upgrade script for Ubuntu 18:

```
$ wget https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/7.x/7.1.7.1/ubuntu18/tars/ambari/ambari-metrics-pre-upgrade-7.1.7.1.14.tar.gz
```

Execute the following command to download the pre-upgrade script for SLES 12:

```
$ wget https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/7.x/7.1.7.1/sles12/tars/ambari/ambari-metrics-pre-upgrade-7.1.7.1.14.tar.gz
```

To find all the download links in a centralised location, see [Software download matrix](#)

b) Extract and run the pre-upgrade script

1. Extract the pre-upgrade tarball: `$ tar xvzf ambari-metrics-pre-upgrade-<version>.tar.gz`
2. Change to the ambari-metrics-pre-upgrade-<version> directory. `$ cd ambari-metrics-pre-upgrade-<version>`
3. Make the pre-upgrade script executable: `$ chmod +x ams-pre-upgrade-2.6-to-7.x-<version>.sh`
4. Run the pre-upgrade script on all Metrics Collector hosts: `$ ./ams-pre-upgrade-2.6-to-7.x-<version>.sh`

c) Stop the service. From Ambari Web, browse to Services > Ambari Metrics and select Stop from the Service Actions menu.

4. If you are running Log Search in your cluster, stop the service. From Ambari Web, browse to Services > Log Search and select Stop from the Service Actions menu. Then, select Turn on Maintenance Mode from the Service Actions menu.
5. Stop the Ambari Server. On the host running Ambari Server: `ambari-server stop`
6. Stop all Ambari Agents. On each host in your cluster running an Ambari Agent: `ambari-agent stop`.
7. Back up your existing Ambari Server database. For example, to back up the default, embedded postgres db for Ambari:

```
mkdir /root/backups/
pg_dump -U ambari ambari > /root/backups/ambari-before-upgrade.sql
```

The default password is: bigdata.



**Note:** You must take a back up of the Ambari database if you have hosted the Ambari database on any other supported databases like MySQL or MariaDB.

8. Upgrade your Ambari Server database to a current supported version, only if your Ambari Server database version is not supported on the target Ambari version. To verify current supported versions, see the *Cloudera Support*

*Matrix*. For example, if you are running postgres v9.4, you must upgrade postgres to v9.6 or v10.2. If you are running MariaDB 10.2, upgrading the Ambari Server database is not required.

See the documentation for your database type and version for specific database upgrade instructions.

9. Retrieve the new Ambari repo and replace the old repository file with the new repository file on all hosts in your cluster.



**Important:** Check your current directory before you download the new repository file to make sure that there are no previous versions of the `ambaridc.repo` file. If you do not, and a previous version exists, the new download is saved with a numeric extension, such as `ambari.repo.1`. Make sure that you copy the new version.

Select the repository appropriate for your environment from the following list:

- For RHEL or CentOS:

```
wget -nv https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/7.x/7.1.7.1/centos7/ambari.repo -O /etc/yum.repos.d/ambari.repo
```

- For Ubuntu:

```
wget -nv https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/7.x/7.1.7.1/ubuntu18/ambari.list -O /etc/apt/sources.list.d/ambari.list
```

- For Sles:

```
wget -nv https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/7.x/7.1.7.1/sles12/ambari.repo -O /etc/zypp/repos.d/ambari.repo
```

10. Upgrade Ambari Server. On the host running Ambari Server:

For RHEL or CentOS:

```
yum clean all
```

```
yum info ambari-server
```

In the information output, visually validate that there is an available version containing "7.1.7.1"

```
yum upgrade ambari-server
```

For Ubuntu 18:

```
apt-get clean all
```

Before upgrading Ambari Server, you must update the username and password in the `ambari.list` file. Run the following command:

```
vi /etc/apt/sources.list.d/ambari.list
```

For example, the output displays the following:

```
#VERSION_NUMBER=7.1.7.1-14
#json.url = https://archive.cloudera.com/p/HDP/hdp_urlinfo.json
deb https://archive.cloudera.com/p/ambaridc/7.x/7.1.7.1/ubuntu18 Ambari main
apt-get update
```



```
apt-cache show ambari-server | grep Version
```

In the information output, visually validate that there is an available version containing "7.1.7.1"

```
apt-get upgrade ambari-server
```

For SLES 12:

```
zypper clean
```

Before upgrading the Ambari Server, you must create, place, and set the credentials for zypper. Run the following command:

```
vi /etc/zypp/repos.d/ambari.repo
```

Add the following postfix to the end of BaseURL and gpgkey: ?credentials=ambari.cat For example, the output displays the following:

```
#VERSION_NUMBER=7.1.7.1-14
[ambari-7.1.7.1]
#json.url = https://archive.cloudera.com/p/HDP/hdp_urlinfo.json
name=ambari Version - ambari-7.1.7.1
baseurl=https://archive.cloudera.com/p/ambaridc/7.x/7.1.7.1/sles12
gpgcheck=0
gpgkey=https://archive.cloudera.com/p/ambaridc/7.x/7.1.7.1/sles12/RPM-GPG-KEY/RPM-GPG-KEY-Jenkins
enabled=1
priority=1
```

Create the following file with your payroll credentials (replace <username> with your payroll username and <password> accordingly): /etc/zypp/credentials.d/ambari.cat

```
username=<username>
password=<password>
```

```
zypper info ambari-server
```

In the information output, visually validate that there is an available version containing "7.1.7.1"

```
Zypper up ambari-server
```

**11.** Check for upgrade success by noting progress during the Ambari Server installation process you started in Step 10.

- As the process runs, the console displays output similar to the following:  
Setting up Upgrade Process Resolving Dependencies Running transaction check
- If the upgrade fails, the console displays output similar to the following:  
Setting up Upgrade Process No Packages marked for Update
- A successful upgrade displays output similar to the following:  
Updated: ambari-server.noarch 0:7.1.7.1. Complete!

**12. Upgrade all Ambari Agents.** On each host in your cluster running an Ambari Agent:

- For RHEL or CentOS: Before upgrading Ambari Agent, you must update the username and password in the `ambaridc.repo` file. Run the following command:

```
vi /etc/yum.repos.d/ambaridc.repo
```

For example, the output displays the following:

```
#VERSION_NUMBER=7.1.7.1-14
[ambari-7.1.7.1]
#json.url = https://archive.cloudera.com/p/HDP/hdp_urlinfo.json
name=ambari Version - ambari-7.1.7.1
baseurl=https://archive.cloudera.com/p/ambaridc/7.x/7.1.7.1/centos7
gpgcheck=1
gpgkey=https://archive.cloudera.com/p/ambaridc/7.x/7.1.7.1/centos7/RPM-GPG-KEY/RPM-GPG-KEY-Jenkins
enabled=1
priority=1
```

```
yum upgrade ambari-agent
```

- For Ubuntu:

```
apt-get update
```

```
apt-get upgrade ambari-agent
```

- For SLES: Copy the following files from the server host to each host with the same path:

```
/etc/zypp/credentials.d/ambari.cat
/etc/zypp/repos.d/ambari.repo
```

Then run the following command:

```
zypper up ambari-agent
```

**13. After the upgrade process completes, check each host to make sure the new files have been installed:**

For CentOS 7.7:

```
rpm -qa | grep ambari-agent
```

For Ubuntu 18:

```
dpkg -l ambari-agent
```

**14. Upgrade Ambari Server database schema.** On the host running Ambari Server: `ambari-server upgrade`. When the Ambari Server database schema has been upgraded, you should see command output like this: Ambari Server ‘upgrade’ completed successfully.**15. Start the Ambari Server.** On the host running Ambari Server: `ambari-server start`**16. Start all Ambari Agents.** On each host in your cluster running an Ambari Agent: `ambari-agent start`**17. Open Ambari Web UI.** Point your browser to the Ambari Web UI:

- When Ambari Server is configured for HTTPS: `https://<your.ambari.server>:8443`
- When Ambari Server is configured for HTTP: `http://<your.ambari.server>:8080`

where `<your.ambari.server>` is the name of your ambari server host. For example, `c7401.ambari.apache.org`.



**Important:** Refresh your browser so that it loads the new version of the Ambari Web code. If you have problems, clear your browser cache manually, then restart Ambari Server.

18. Log in using the Ambari administrator credentials that you have set up. For example, the default name/password is admin/admin. You will see a Restart indicator next to each service after upgrading. Ambari upgrade has modified the configuration properties of your cluster based on the new configuration types and properties being made available for each service with this release of Ambari. Review these changes by comparing the previous configuration with the latest version created by "ambari-upgrade".



**Caution:** Do not manually restart these services unless future steps in the upgrade guide prompt you to do so. Manually restarting these services may significantly disrupt your upgrade. Ambari will restart each service automatically during the HDP upgrade.

19. If you have configured Ambari to authenticate against an external LDAP or Active Directory, you must re-run `ambari-server setup-ldap`.
20. If you have configured your cluster for Hive, Ranger or Oozie with an external database (Oracle, MySQL or PostgreSQL), you must run the following command to get the JDBC driver jar file in place: `ambari-server setup --jdbc-db` and `--jdbc-driver`.  
A sample command to setup MySQL connector: `ambari-server setup --jdbc-db=mysql --jdbc-driver=/path/to/mysql/mysql-connector-java.jar_`

## Upgrade to Ambari 7.1.6.0

After reviewing the information related to Ambari UI and the Quick Links, and backing up the Ambari Server configuration file, perform the recommended steps for upgrading Ambari.

### Procedure

1. Turn off Auto Restart from Ambari web UI by browsing to Admin > Service Auto Start. Set Auto-Start Services to Disabled. Click Save.
2. If you are running SmartSense in your cluster, stop the service. From Ambari Web, browse to Services > SmartSense and select Stop from the Service Actions menu. Then, select Turn on Maintenance Mode from the Service Actions menu.

3. If you are running Ambari Metrics in your cluster, put it in Maintenance Mode. Don't stop the service yet. From Ambari Web, browse to Services > Ambari Metrics and select Turn on Maintenance Mode from the Service Actions menu.

a) Download the pre-upgrade script

1. SSH into the host running the Metrics Collector. Note, that in case Ambari Metrics System (AMS) is running in High Availability mode, there are multiple hosts running the Metrics Collector. These steps must be done on each Metrics Collector host.
2. Become the AMS service user. For example, run the following su command on Linux: `$ sudo su - ams`
3. Change to the /tmp directory.
4. Execute the following command to download the pre-upgrade script for CentOS 7:

```
$ wget https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/7.x/7.1.6.0/centos7/tars/ambari/ambari-metrics-pre-upgrade-7.1.6.0.38.tar.gz
```

Execute the following command to download the pre-upgrade script for Ubuntu 18:

```
$ wget https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/7.x/7.1.6.0/ubuntu18/tars/ambari/ambari-metrics-pre-upgrade-7.1.6.0.38.tar.gz
```

Execute the following command to download the pre-upgrade script for SLES 12:

```
$ wget https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/7.x/7.1.6.0/sles12/tars/ambari/ambari-metrics-pre-upgrade-7.1.6.0.38.tar.gz
```

To find all the download links in a centralised location, see [Software download matrix](#)

b) Extract and run the pre-upgrade script

1. Extract the pre-upgrade tarball: `$ tar xvzf ambari-metrics-pre-upgrade-<version>.tar.gz`
2. Change to the ambari-metrics-pre-upgrade-<version> directory. `$ cd ambari-metrics-pre-upgrade-<version>`
3. Make the pre-upgrade script executable: `$ chmod +x ams-pre-upgrade-2.6-to-7.x-<version>.sh`
4. Run the pre-upgrade script on all Metrics Collector hosts: `$ ./ams-pre-upgrade-2.6-to-7.x-<version>.sh`

c) Stop the service. From Ambari Web, browse to Services > Ambari Metrics and select Stop from the Service Actions menu.

4. If you are running Log Search in your cluster, stop the service. From Ambari Web, browse to Services > Log Search and select Stop from the Service Actions menu. Then, select Turn on Maintenance Mode from the Service Actions menu.
5. Stop the Ambari Server. On the host running Ambari Server: `ambari-server stop`
6. Stop all Ambari Agents. On each host in your cluster running an Ambari Agent: `ambari-agent stop`.
7. Back up your existing Ambari Server database. For example, to back up the default, embedded postgres db for Ambari:

```
mkdir /root/backups/
pg_dump -U ambari ambari > /root/backups/ambari-before-upgrade.sql
```

The default password is: bigdata.



**Note:** You must take a back up of the Ambari database if you have hosted the Ambari database on any other supported databases like MySQL or MariaDB.

8. Upgrade your Ambari Server database to a current supported version, only if your Ambari Server database version is not supported on the target Ambari version. To verify current supported versions, see the *Cloudera Support*

*Matrix*. For example, if you are running postgres v9.4, you must upgrade postgres to v9.6 or v10.2. If you are running MariaDB 10.2, upgrading the Ambari Server database is not required.

See the documentation for your database type and version for specific database upgrade instructions.

- Retrieve the new Ambari repo and replace the old repository file with the new repository file on all hosts in your cluster.



**Important:** Check your current directory before you download the new repository file to make sure that there are no previous versions of the `ambaridc.repo` file. If you do not, and a previous version exists, the new download is saved with a numeric extension, such as `ambaridc.repo.1`. Make sure that you copy the new version.

Select the repository appropriate for your environment from the following list:

- For RHEL or CentOS:

```
wget -nv https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/7.x/7.1.6.0/centos7/ambaridc.repo -O /etc/yum.repos.d/ambaridc.repo
```

- For Ubuntu:

```
wget -nv https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/7.x/7.1.6.0/ubuntu18/ambaridc.list -O /etc/apt/sources.list.d/ambaridc.list
```

- For Sles:

```
wget -nv https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/7.x/7.1.6.0/sles12/ambaridc.repo -O /etc/zypp/repos.d/ambari.repo
```

- Upgrade Ambari Server. On the host running Ambari Server:

For RHEL or CentOS:

```
yum clean all
```

```
yum info ambari-server
```

In the information output, visually validate that there is an available version containing "7.1.6.0"

```
yum upgrade ambari-server
```

For Ubuntu 18:

```
apt-get clean all
```

Before upgrading Ambari Server, you must update the username and password in the `ambari.list` file. Run the following command:

```
vi /etc/apt/sources.list.d/ambaridc.list
```

For example, the output displays the following:

```
#VERSION_NUMBER=7.1.6.0-38
#json.url = http://public-repo-1.hortonworks.com/HDP/hdp_urlinfo.json
deb https://archive.cloudera.com/p/ambaridc/7.x/7.1.6.0/ubuntu18/ Ambari
main
apt-get update
```

```
apt-cache show ambari-server | grep Version
```

In the information output, visually validate that there is an available version containing "7.1.6.0"

```
apt-get upgrade ambari-server
```

For SLES 12:

```
zypper clean
```

Before upgrading the Ambari Server, you must create, place, and set the credentials for zypper. Run the following command:

```
vi /etc/zypp/repos.d/ambari.repo
```

Add the following postfix to the end of BaseURL and gpgkey: `?credentials=ambari.cat` For example, the output displays the following:

```
#VERSION_NUMBER=7.1.6.0-38
[ambari-7.1.6.0]
#json.url = http://public-repo-1.hortonworks.com/HDP/hdp_urlinfo.json
name=ambari Version - ambari-7.1.6.0
baseurl=https://archive.cloudera.com/p/ambaridc/7.x/7.1.6.0/sles12/
gpgcheck=0
gpgkey=https://archive.cloudera.com/p/ambaridc/7.x/7.1.6.0/sles12/RPM-GPG-KEY/RPM-GPG-KEY-Jenkins
enabled=1
priority=1
```

Create the following file with your paywall credentials (replace `<username>` with your paywall username and `<password>` accordingly): `/etc/zypp/credentials.d/ambari.cat`

```
username=<username>
password=<password>
```

```
zypper info ambari-server
```

In the information output, visually validate that there is an available version containing "7.1.6.0"

```
Zypper up ambari-server
```

**11.** Check for upgrade success by noting progress during the Ambari Server installation process you started in Step 10.

- As the process runs, the console displays output similar to the following:  
Setting up Upgrade Process Resolving Dependencies Running transaction check
- If the upgrade fails, the console displays output similar to the following:  
Setting up Upgrade Process No Packages marked for Update
- A successful upgrade displays output similar to the following:  
Updated: ambari-server.noarch 0:7.1.6.0. Complete!

**12. Upgrade all Ambari Agents.** On each host in your cluster running an Ambari Agent:

- For RHEL or CentOS: Before upgrading Ambari Agent, you must update the username and password in the `ambaridc.repo` file. Run the following command:

```
vi /etc/yum.repos.d/ambaridc.repo
```

For example, the output displays the following:

```
#VERSION_NUMBER=7.1.6.0-38
[ambari-7.1.6.0]
#json.url = http://public-repo-1.hortonworks.com/HDP/hdp_urlinfo.json
name=ambari Version - ambari-7.1.6.0
baseurl=https://archive.cloudera.com/p/ambaridc/7.x/7.1.6.0/centos7/
gpgcheck=1
gpgkey=https://archive.cloudera.com/p/ambaridc/7.x/7.1.6.0/centos7/RPM-
GPG-KEY/RPM-GPG-KEY-Jenkins
enabled=1
priority=1
```

```
yum upgrade ambari-agent
```

- For Ubuntu:

```
apt-get update
```

```
apt-get upgrade ambari-agent
```

- For SLES: Copy the following files from the server host to each host with the same path:

```
/etc/zypp/credentials.d/ambari.cat
/etc/zypp/repos.d/ambari.repo
```

Then run the following command:

```
zypper up ambari-agent
```

**13. After the upgrade process completes, check each host to make sure the new files have been installed:**

For CentOS 7.7:

```
rpm -qa | grep ambari-agent
```

For Ubuntu 18:

```
dpkg -l ambari-agent
```

**14. Upgrade Ambari Server database schema.** On the host running Ambari Server: `ambari-server upgrade`. When the Ambari Server database schema has been upgraded, you should see command output like this: Ambari Server ‘upgrade’ completed successfully.**15. Start the Ambari Server.** On the host running Ambari Server: `ambari-server start`**16. Start all Ambari Agents.** On each host in your cluster running an Ambari Agent: `ambari-agent start`**17. Open Ambari Web UI.** Point your browser to the Ambari Web UI:

- When Ambari Server is configured for HTTPS: `https://<your.ambari.server>:8443`
- When Ambari Server is configured for HTTP: `http://<your.ambari.server>:8080`

where `<your.ambari.server>` is the name of your ambari server host. For example, `c7401.ambari.apache.org`.



**Important:** Refresh your browser so that it loads the new version of the Ambari Web code. If you have problems, clear your browser cache manually, then restart Ambari Server.

18. Log in using the Ambari administrator credentials that you have set up. For example, the default name/password is admin/admin. You will see a Restart indicator next to each service after upgrading. Ambari upgrade has modified the configuration properties of your cluster based on the new configuration types and properties being made available for each service with this release of Ambari. Review these changes by comparing the previous configuration with the latest version created by "ambari-upgrade".



**Caution:** Do not manually restart these services unless future steps in the upgrade guide prompt you to do so. Manually restarting these services may significantly disrupt your upgrade. Ambari will restart each service automatically during the HDP upgrade.

19. If you have configured Ambari to authenticate against an external LDAP or Active Directory, you must re-run `ambari-server setup-ldap`.
20. If you have configured your cluster for Hive, Ranger or Oozie with an external database (Oracle, MySQL or PostgreSQL), you must run the following command to get the JDBC driver jar file in place: `ambari-server setup --jdbc-db` and `--jdbc-driver`.  
A sample command to setup MySQL connector: `ambari-server setup --jdbc-db=mysql --jdbc-driver=/path/to/mysql/mysql-connector-java.jar`

## Upgrade to Ambari 7.1.4.0

After reviewing the information related to behavioral changes, Ambari UI and the Quick Links, and backing up the Ambari Server configuration file, perform the recommended steps for upgrading Ambari.

### Procedure

1. Turn off Auto Restart from Ambari web UI by browsing to Admin > Service Auto Start. Set Auto-Start Services to Disabled. Click Save.
2. If you are running SmartSense in your cluster, stop the service. From Ambari Web, browse to Services > SmartSense and select Stop from the Service Actions menu. Then, select Turn on Maintenance Mode from the Service Actions menu.



3. If you are running Ambari Metrics in your cluster, put it in Maintenance Mode. Don't stop the service yet. From Ambari Web, browse to Services > Ambari Metrics and select Turn on Maintenance Mode from the Service Actions menu.

a) Download the pre-upgrade script

1. SSH into the host running the Metrics Collector. Note, that in case Ambari Metrics System (AMS) is running in High Availability mode, there are multiple hosts running the Metrics Collector. These steps must be done on each Metrics Collector host.
2. Become the AMS service user. For example, run the following su command on Linux: `$ sudo su - ams`
3. Change to the /tmp directory.
4. Execute the following command to download the pre-upgrade script for CentOS 7:

```
$ wget https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/centos7/7.x/updates/7.1.4.0/tars/ambari/ambari-metrics-pre-upgrade-7.1.4.0.25.tar.gz
```

Execute the following command to download the pre-upgrade script for Ubuntu 18:

```
$ wget https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/ubuntu18/7.x/updates/7.1.4.0/tars/ambari/ambari-metrics-pre-upgrade-7.1.4.0.25.tar.gz
```

Execute the following command to download the pre-upgrade script for SLES 12:

```
$ wget https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/sles12/7.x/updates/7.1.4.0/tars/ambari/ambari-metrics-pre-upgrade-7.1.4.0.25.tar.gz
```

To find all the download links in a centralised location, see [Software download matrix](#)

b) Extract and run the pre-upgrade script

1. Extract the pre-upgrade tarball: `$ tar xvzf ambari-metrics-pre-upgrade-<version>.tar.gz`
2. Change to the ambari-metrics-pre-upgrade-<version> directory. `$ cd ambari-metrics-pre-upgrade-<version>`
3. Make the pre-upgrade script executable: `$ chmod +x ambari-metrics-pre-upgrade-2.6-to-7.x-<version>.sh`
4. Run the pre-upgrade script on all Metrics Collector hosts: `$ ./ambari-metrics-pre-upgrade-2.6-to-7.x-<version>.sh`

Ambari Metrics Collector will shut down automatically.

- c) Stop the service. From Ambari Web, browse to Services > Ambari Metrics and select Stop from the Service Actions menu.
4. If you are running Log Search in your cluster, stop the service. From Ambari Web, browse to Services > Log Search and select Stop from the Service Actions menu. Then, select Turn on Maintenance Mode from the Service Actions menu.
5. Stop the Ambari Server. On the host running Ambari Server: `ambari-server stop`
6. Stop all Ambari Agents. On each host in your cluster running an Ambari Agent: `ambari-agent stop`.
7. Back up your existing Ambari Server database. For example, to back up the default, embedded postgres db for Ambari:

```
mkdir /root/backups/
pg_dump -U ambari ambari > /root/backups/ambari-before-upgrade.sql
```

The default password is: bigdata.



**Note:** You must take a back up of the Ambari database if you have hosted the Ambari database on any other supported databases like MySQL or MariaDB.

8. Upgrade your Ambari Server database to a current supported version, only if your Ambari Server database version is not supported on the target Ambari version. To verify current supported versions, see [Cloudera Support Matrix](#). For example, if you are running postgres v9.4, you must upgrade postgres to v9.6 or v10.2. If you are running MariaDB 10.2, upgrading the Ambari Server database is not required.

See the documentation for your database type and version for specific database upgrade instructions.

9. Retrieve the new Ambari repo and replace the old repository file with the new repository file on all hosts in your cluster.



**Important:** Check your current directory before you download the new repository file to make sure that there are no previous versions of the ambaridc.repo file. If you do not, and a previous version exists, the new download is saved with a numeric extension, such as ambaridc.repo.1. Make sure that you copy the new version.

Select the repository appropriate for your environment from the following list:

- For RHEL or CentOS:

```
wget -nv https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/centos7/7.x/updates/7.1.4.0/ambaridc.repo -O /etc/yum.repos.d/ambaridc.repo
```

- For Ubuntu:

```
wget -nv https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/ubuntu18/7.x/updates/7.1.4.0/ambaridc.list -O /etc/apt/sources.list.d/ambaridc.list
```

- For Sles:

```
wget -nv https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/sles12/7.x/updates/7.1.4.0/ambaridc.repo -O /etc/zypp/repos.d/ambari.repo
```

10. Upgrade Ambari Server. On the host running Ambari Server:

For RHEL or CentOS:

```
yum clean all
```

```
yum info ambari-server
```

In the information output, visually validate that there is an available version containing "7.1.4.0"

```
yum upgrade ambari-server
```

For Ubuntu 18:

```
apt-get clean all
```

Before upgrading Ambari Server, you must update the username and password in the ambari.list file. Run the following command:

```
vi /etc/apt/sources.list.d/ambaridc.list
```

For example, the output displays the following:

```
#VERSION_NUMBER=7.1.4.0-25
#json.url = http://public-repo-1.hortonworks.com/HDP/hdp_urlinfo.json
deb https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/ubuntu18/7.x/updates/7.1.4.0 Ambari main
```

```
apt-get update
apt-cache show ambari-server | grep Version
```

In the information output, visually validate that there is an available version containing "7.1.4.0"

```
apt-get upgrade ambari-server
```

For SLES 12:

```
zypper clean
```

Before upgrading the Ambari Server, you must create, place, and set the credentials for zypper. Run the following command:

```
vi /etc/zypp/repos.d/ambari.repo
```

Add the following postfix to the end of BaseURL and gpgkey: ?credentials=ambari.cat For example, the output displays the following:

```
#VERSION_NUMBER=7.1.4.0-25
[ambari-7.1.4.0]
#json.url = http://public-repo-1.hortonworks.com/HDP/hdp_urlinfo.json
name=ambari Version - ambari-7.1.4.0
baseurl=https://archive.cloudera.com/p/ambaridc/sles12/7.x/updates/7.1.4.0?credentials=ambari.cat
gpgcheck=1
gpgkey=https://archive.cloudera.com/p/ambaridc/sles12/7.x/updates/7.1.4.0/RPM-GPG-KEY/RPM-GPG-KEY-Jenkins?credentials=ambari.cat
enabled=1
priority=1
```

Create the following file with your payroll credentials (replace <username> with your payroll username and <password> accordingly): /etc/zypp/credentials.d/ambari.cat

```
username=<username>
password=<password>
```

```
zypper info ambari-server
```

In the information output, visually validate that there is an available version containing "7.1.4.0"

```
Zypper up ambari-server
```

**11.** Check for upgrade success by noting progress during the Ambari Server installation process you started in Step 10.

- As the process runs, the console displays output similar to the following:  
Setting up Upgrade Process Resolving Dependencies Running transaction check
- If the upgrade fails, the console displays output similar to the following:  
Setting up Upgrade Process No Packages marked for Update
- A successful upgrade displays output similar to the following:  
Updated: ambari-server.noarch 0:7.1.4.0. Complete!

**12. Upgrade all Ambari Agents. On each host in your cluster running an Ambari Agent:**

- For RHEL or CentOS: Before upgrading Ambari Agent, you must update the username and password in the `ambaridc.repo` file. Run the following command:

```
vi /etc/yum.repos.d/ambaridc.repo
```

For example, the output displays the following:

```
#VERSION_NUMBER=7.1.4.0-25
[ambari-7.1.4.0]
#json.url = http://public-repo-1.hortonworks.com/HDP/hdp_urlinfo.json
name=ambari Version - ambari-7.1.4.0
baseurl=https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/
p/ambaridc/centos7/7.x/updates/7.1.4.0
gpgcheck=1
gpgkey=https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/
ambaridc/centos7/7.x/updates/7.1.4.0/RPM-GPG-KEY/RPM-GPG-KEY-Jenkins
enabled=1
priority=1
```

```
yum upgrade ambari-agent
```

- For Ubuntu:

```
apt-get update
```

```
apt-get upgrade ambari-agent
```

- For SLES: Copy the following files from the server host to each host with the same path:

```
/etc/zypp/credentials.d/ambari.cat
/etc/zypp/repos.d/ambari.repo
```

Then run the following command:

```
zypper up ambari-agent
```

**13. After the upgrade process completes, check each host to make sure the new files have been installed:**

For CentOS 7.7:

```
rpm -qa | grep ambari-agent
```

For Ubuntu 18:

```
dpkg -l ambari-agent
```

**14. Upgrade Ambari Server database schema. On the host running Ambari Server: `ambari-server upgrade`. When the Ambari Server database schema has been upgraded, you should see command output like this: Ambari Server ‘u pgrade’ completed successfully.****15. Start the Ambari server from the host running using the following command:**

```
ambari-server start
```

**16. Start all Ambari Agents on each host in your cluster running an Ambari Agent using the following command:**

```
ambari-agent start
```

17. Open Ambari Web UI. Point your browser to the Ambari Web UI:

- When Ambari Server is configured for HTTPS: `https://<your.ambari.server>:8443`
- When Ambari Server is configured for HTTP: `http://<your.ambari.server>:8080`

where `<your.ambari.server>` is the name of your ambari server host. For example, `c7401.ambari.apache.org`.



**Important:** Refresh your browser so that it loads the new version of the Ambari Web code. If you have problems, clear your browser cache manually, then restart Ambari Server.

18. Log in using the Ambari administrator credentials that you have set up. For example, the default name/password is `admin/admin`. You will see a Restart indicator next to each service after upgrading. Ambari upgrade has modified the configuration properties of your cluster based on the new configuration types and properties being made available for each service with this release of Ambari. Review these changes by comparing the previous configuration with the latest version created by "ambari-upgrade".



**Caution:** Do not manually restart these services unless future steps in the upgrade guide prompt you to do so. Manually restarting these services may significantly disrupt your upgrade. Ambari will restart each service automatically during the HDP upgrade.

19. If you have configured Ambari to authenticate against an external LDAP or Active Directory, you must re-run `ambari-server setup-ldap`.

20. If you have configured your cluster for Hive, Ranger or Oozie with an external database (Oracle, MySQL or PostgreSQL), you must run the following command to get the JDBC driver jar file in place: `ambari-server setup --jdbc-db` and `--jdbc-driver`.

A sample command to setup MySQL connector: `ambari-server setup --jdbc-db=mysql --jdbc-driver=/path/to/mysql/mysql-connector-java.jar_`

Ambari is ready and upgraded to 7.1.x.

## Upgrade to Ambari 7.1.3.0

After reviewing the information related to behavioral changes, Ambari UI and the Quick Links, and backing up the Ambari Server configuration file, perform the recommended steps for upgrading Ambari.

### Procedure

1. Turn off Auto Restart from Ambari web UI by browsing to Admin > Service Auto Start. Set Auto-Start Services to Disabled. Click Save.
2. If you are running SmartSense in your cluster, stop the service. From Ambari Web, browse to Services > SmartSense and select Stop from the Service Actions menu. Then, select Turn on Maintenance Mode from the Service Actions menu.

3. If you are running Ambari Metrics in your cluster, put it in Maintenance Mode. Don't stop the service yet. From Ambari Web, browse to Services > Ambari Metrics and select Turn on Maintenance Mode from the Service Actions menu.

a) Download the pre-upgrade script

1. SSH into the host running the Metrics Collector. Note, that in case Ambari Metrics System (AMS) is running in High Availability mode, there are multiple hosts running the Metrics Collector. These steps must be done on each Metrics Collector host.
2. Become the AMS service user. For example, run the following su command on Linux: `$ sudo su - ams`
3. Change to the /tmp directory.
4. Execute the following command to download the pre-upgrade script for CentOS 7:

```
$ wget https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/centos7/7.x/updates/7.1.3.0/tars/ambari/ambari-metrics-pre-upgrade-7.1.3.0.17.tar.gz
```

Execute the following command to download the pre-upgrade script for Ubuntu 18:

```
$ wget https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/ubuntu18/7.x/updates/7.1.3.0/tars/ambari/ambari-metrics-pre-upgrade-7.1.3.0.17.tar.gz
```

To find all the download links in a centralised location, see [Software download matrix](#)

b) Extract and run the pre-upgrade script

1. Extract the pre-upgrade tarball: `$ tar xvzf ambari-metrics-pre-upgrade-<version>.tar.gz`
2. Change to the ambari-metrics-pre-upgrade-<version> directory. `$ cd ambari-metrics-pre-upgrade-<version>`
3. Make the pre-upgrade script executable: `$ chmod +x ambari-metrics-pre-upgrade-2.6-to-7.x-<version>.sh`
4. Run the pre-upgrade script on all Metrics Collector hosts: `$ ./ambari-metrics-pre-upgrade-2.6-to-7.x-<version>.sh`

Ambari Metrics Collector will shut down automatically.

c) Stop the service. From Ambari Web, browse to Services > Ambari Metrics and select Stop from the Service Actions menu.

4. If you are running Log Search in your cluster, stop the service. From Ambari Web, browse to Services > Log Search and select Stop from the Service Actions menu. Then, select Turn on Maintenance Mode from the Service Actions menu.
5. Stop the Ambari Server. On the host running Ambari Server: `ambari-server stop`
6. Stop all Ambari Agents. On each host in your cluster running an Ambari Agent: `ambari-agent stop`.
7. Back up your existing Ambari Server database. For example, to back up the default, embedded postgres db for Ambari:

```
mkdir /root/backups/
pg_dump -U ambari ambari > /root/backups/ambari-before-upgrade.sql
```

The default password is: bigdata.



**Note:** You must take a back up of the Ambari database if you have hosted the Ambari database on any other supported databases like MySQL or MariaDB.

8. Upgrade your Ambari Server database to a current supported version, only if your Ambari Server database version is not supported on the target Ambari version. To verify current supported versions, see the [Cloudera Support Matrix](#). For example, if you are running postgres v9.4, you must upgrade postgres to v9.6 or v10.2. If you are running MariaDB 10.2, upgrading the Ambari Server database is not required.

See the documentation for your database type and version for specific database upgrade instructions.

- Retrieve the new Ambari repo and replace the old repository file with the new repository file on all hosts in your cluster.



**Important:** Check your current directory before you download the new repository file to make sure that there are no previous versions of the `ambaridc.repo` file. If you do not, and a previous version exists, the new download is saved with a numeric extension, such as `ambaridc.repo.1`. Make sure that you copy the new version.

Select the repository appropriate for your environment from the following list:

- For RHEL or CentOS:

```
wget -nv https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/
p/ambaridc/centos7/7.x/updates/7.1.3.0/
ambaridc.repo -O /etc/yum.repos.d/ambaridc.repo
```

- For Ubuntu:

```
wget -nv https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com
/p/ambaridc/ubuntu18/7.x/updates/7.1.3.0/
ambaridc.list -O /etc/apt/sources.list.d/ambaridc.list
```

- Upgrade Ambari Server. On the host running Ambari Server:

For RHEL or CentOS:

```
yum clean all
```

```
yum info ambari-server
```

In the information output, visually validate that there is an available version containing "7.1.3.0"

```
yum upgrade ambari-server
```

For Ubuntu 18:

```
apt-get clean all
```

Before upgrading Ambari Server, you must update the username and password in the `ambari.list` file. Run the following command:

```
vi /etc/apt/sources.list.d/ambaridc.list
```

For example, the output displays the following:

```
#VERSION_NUMBER=7.1.3.0-17
#json.url = http://public-repo-1.hortonworks.com/HDP/hdp_urlinfo.json
deb https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/am
baridc/ubuntu18/7.x/updates/7.1.3.0 Ambari main
apt-get update
apt-cache show ambari-server | grep Version
```

In the information output, visually validate that there is an available version containing "7.1.3.0"

```
apt-get install ambari-server
```

- 11.** Check for upgrade success by noting progress during the Ambari Server installation process you started in Step 10.
- As the process runs, the console displays output similar to the following:  
Setting up Upgrade Process Resolving Dependencies Running transaction check
  - If the upgrade fails, the console displays output similar to the following:  
Setting up Upgrade Process No Packages marked for Update
  - A successful upgrade displays output similar to the following:

```
Updated: ambari-server.noarch 0:7.1.3.0. Complete!
```

- 12.** Upgrade all Ambari Agents. On each host in your cluster running an Ambari Agent:

- For RHEL or CentOS: Before upgrading Ambari Agent, you must update the username and password in the `ambaridc.repo` file. Run the following command:

```
vi /etc/yum.repos.d/ambaridc.repo
```

For example, the output displays the following:

```
#VERSION_NUMBER=7.1.3.0-17
[ambari-7.1.3.0]
#json.url = http://public-repo-1.hortonworks.com/HDP/hdp_urlinfo.json
name=ambari Version - ambari-7.1.3.0
baseurl=https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/centos7/7.x/updates/7.1.3.0
gpgcheck=1
gpgkey=https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/centos7/7.x/updates/7.1.3.0/RPM-GPG-KEY/RPM-GPG-KEY-Jenkins
enabled=1
priority=1
```

```
yum upgrade ambari-agent
```

- For Ubuntu:

```
apt-get update
apt-get install ambari-agent
```

- 13.** After the upgrade process completes, check each host to make sure the new files have been installed:

For CentOS 7.7:

```
rpm -qa | grep ambari-agent
```

For Ubuntu 18:

```
dpkg -l ambari-agent
```

- 14.** Upgrade Ambari Server database schema. On the host running Ambari Server: `ambari-server upgrade`. When the Ambari Server database schema has been upgraded, you should see command output like this: Ambari Server ‘`u pgrade`’ completed successfully.
- 15.** Start the Ambari server from the host running using the following command:
- ```
ambari-server start
```
- 16.** Start all Ambari Agents on each host in your cluster running an Ambari Agent using the following command:
- ```
ambari-agent start
```



17. Open Ambari Web UI. Point your browser to the Ambari Web UI:

- When Ambari Server is configured for HTTPS: `https://<your.ambari.server>:8443`
- When Ambari Server is configured for HTTP: `http://<your.ambari.server>:8080`

where `<your.ambari.server>` is the name of your ambari server host. For example, `c7401.ambari.apache.org`.



**Important:** Refresh your browser so that it loads the new version of the Ambari Web code. If you have problems, clear your browser cache manually, then restart Ambari Server.

18. Log in using the Ambari administrator credentials that you have set up. For example, the default name/password is `admin/admin`. You will see a Restart indicator next to each service after upgrading. Ambari upgrade has modified the configuration properties of your cluster based on the new configuration types and properties being made available for each service with this release of Ambari. Review these changes by comparing the previous configuration with the latest version created by "ambari-upgrade".



**Caution:** Do not manually restart these services unless future steps in the upgrade guide prompt you to do so. Manually restarting these services may significantly disrupt your upgrade. Ambari will restart each service automatically during the HDP upgrade.

19. If you have configured Ambari to authenticate against an external LDAP or Active Directory, you must re-run `ambari-server setup-ldap`.

20. If you have configured your cluster for Hive, Ranger or Oozie with an external database (Oracle, MySQL or PostgreSQL), you must run the following command to get the JDBC driver jar file in place: `ambari-server setup --jdbc-db` and `--jdbc-driver`.

A sample command to setup MySQL connector: `ambari-server setup --jdbc-db=mysql --jdbc-driver=/path/to/mysql/mysql-connector-java.jar_`

Ambari is ready and upgraded to 7.1.x.

## Upgrade to Ambari 7.1.2.0

After reviewing the information related to behavioral changes, Ambari UI and the Quick Links, and backing up the Ambari Server configuration file, perform the recommended steps for upgrading Ambari.

### Procedure

1. Turn off Auto Restart from Ambari web UI by browsing to Admin > Service Auto Start. Set Auto-Start Services to Disabled. Click Save.
2. If you are running SmartSense in your cluster, stop the service. From Ambari Web, browse to Services > SmartSense and select Stop from the Service Actions menu. Then, select Turn on Maintenance Mode from the Service Actions menu.

3. If you are running Ambari Metrics in your cluster, put it in Maintenance Mode. Don't stop the service yet. From Ambari Web, browse to Services > Ambari Metrics and select Turn on Maintenance Mode from the Service Actions menu.

a) Download the pre-upgrade script

1. SSH into the host running the Metrics Collector. Note, that in case Ambari Metrics System (AMS) is running in High Availability mode, there are multiple hosts running the Metrics Collector. These steps must be done on each Metrics Collector host.
2. Become the AMS service user. For example, run the following su command on Linux: `$ sudo su - ams`
3. Change to the /tmp directory.
4. Execute the following command to download the pre-upgrade script:

```
$ wget https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/centos7/7.x/updates/7.1.2.0/tars/ambari/ambari-metrics-pre-upgrade-7.1.2.0.11.tar.gz
```

To find all the download links in a centralised location, see [Software download matrix](#)

b) Extract and run the pre-upgrade script

1. Extract the pre-upgrade tarball: `$ tar xvzf ambari-metrics-pre-upgrade-<version>.tar.gz`
2. Change to the ambari-metrics-pre-upgrade-<version> directory. `$ cd ambari-metrics-pre-upgrade-<version>`
3. Make the pre-upgrade script executable: `$ chmod +x amms-pre-upgrade-2.6-to-7.x-<version>.sh`
4. Run the pre-upgrade script on all Metrics Collector hosts: `$ ./ams-pre-upgrade-2.6-to-7.x-<version>.sh`

Ambari Metrics Collector will shut down automatically.

- c) Stop the service. From Ambari Web, browse to Services > Ambari Metrics and select Stop from the Service Actions menu.
4. If you are running Log Search in your cluster, stop the service. From Ambari Web, browse to Services > Log Search and select Stop from the Service Actions menu. Then, select Turn on Maintenance Mode from the Service Actions menu.
5. Stop the Ambari Server. On the host running Ambari Server: `ambari-server stop`
6. Stop all Ambari Agents. On each host in your cluster running an Ambari Agent: `ambari-agent stop`.
7. Back up your existing Ambari Server database. For example, to back up the default, embedded postgres db for Ambari:

```
mkdir /root/backups/
pg_dump -U ambari ambari > /root/backups/ambari-before-upgrade.sql
```

The default password is: bigdata.



**Note:** You must take a back up of the Ambari database if you have hosted the Ambari database on any other supported databases like MySQL or MariaDB.

8. Upgrade your Ambari Server database to a current supported version, only if your Ambari Server database version is not supported on the target Ambari version. To verify current supported versions, see the [Cloudera Support Matrix](#). For example, if you are running postgres v9.4, you must upgrade postgres to v9.6 or v10.2. If you are running MariaDB 10.2, upgrading the Ambari Server database is not required.

See the documentation for your database type and version for specific database upgrade instructions.

- Retrieve the new Ambari repo and replace the old repository file with the new repository file on all hosts in your cluster.



**Important:** Check your current directory before you download the new repository file to make sure that there are no previous versions of the `ambaridc.repo` file. If you do not, and a previous version exists, the new download is saved with a numeric extension, such as `ambaridc.repo.1`. Make sure that you copy the new version.

Select the repository appropriate for your environment from the following list:

- For RHEL or CentOS:

```
wget -nv https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/
p/ambaridc/centos7/7.x/updates/7.1.2.0/
                                ambaridc.repo -O /etc/yum.repos.d/ambari
dc.repo
```

- For Ubuntu:

```
wget -nv https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com
/p/ambaridc/ubuntu18/7.x/updates/7.1.2.0/
                                ambaridc.list -O /etc/apt/sources.list
.d/ambaridc.list
```

- Upgrade Ambari Server. On the host running Ambari Server:

For RHEL or CentOS:

```
yum clean all
```

```
yum info ambari-server
```

In the information output, visually validate that there is an available version containing "7.1.2.0"

```
yum upgrade ambari-server
```

For Ubuntu 18:

```
apt-get clean all
```

Before upgrading Ambari Server, you must update the username and password in the `ambari.list` file. Run the following command:

```
vi /etc/apt/sources.list.d/ambaridc.list
```

For example, the output displays the following:

```
#VERSION_NUMBER=7.1.2.0-11
#json.url = http://public-repo-1.hortonworks.com/
HDP/hdp_urlinfo.json
deb https://[***USERNAME***]:[***PASSWORD***]@arch
ive.cloudera.com/p/ambaridc/ubuntu18/7.x/updates/7.1.2.0 Ambari main

apt-get update
```

```
apt-cache show ambari-server | grep Version
```

In the information output, visually validate that there is an available version containing "7.1.2.0"

```
apt-get install ambari-server
```

**11.** Check for upgrade success by noting progress during the Ambari Server installation process you started in Step 10.

- As the process runs, the console displays output similar to the following:  
Setting up Upgrade Process Resolving Dependencies Running transaction check
- If the upgrade fails, the console displays output similar to the following:  
Setting up Upgrade Process No Packages marked for Update
- A successful upgrade displays output similar to the following:  
Updated: ambari-server.noarch 0:7.1.2.0. Complete!

**12.** Upgrade all Ambari Agents. On each host in your cluster running an Ambari Agent:

- For RHEL or CentOS: Before upgrading Ambari Agent, you must update the username and password in the ambaridc.repo file. Run the following command:

```
vi /etc/yum.repos.d/ambaridc.repo
```

For example, the output displays the following:

```
#VERSION_NUMBER=7.1.2.0-11
[ambari-7.1.2.0]
#json.url = http://public-repo-1.hortonworks.com/HDP/hdp_urlinfo.json
name=ambari Version - ambari-7.1.2.0
baseurl=https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/
p/ambaridc/centos7/7.x/updates/7.1.2.0
gpgcheck=1
gpgkey=https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/
ambaridc/centos7/7.x/updates/7.1.2.0/RPM-GPG-KEY/RPM-GPG-KEY-Jenkins
enabled=1
priority=1
```

```
yum upgrade ambari-agent
```

- For Ubuntu:

```
apt-get update
```

```
apt-get install ambari-agent
```

**13.** After the upgrade process completes, check each host to make sure the new files have been installed:

For CentOS 7.7:

```
rpm -qa | grep ambari-agent
```

For Ubuntu 18:

```
dpkg -l ambari-agent
```

**14.** Upgrade Ambari Server database schema. On the host running Ambari Server: `ambari-server upgrade`. When the Ambari Server database schema has been upgraded, you should see command output like this: Ambari Server 'u pgrade' completed successfully.

**15.** Start the Ambari server from the host running using the following command:

```
ambari-server start
```

16. Start all Ambari Agents on each host in your cluster running an Ambari Agent using the following command:

```
ambari-agent start
```

17. Open Ambari Web UI. Point your browser to the Ambari Web UI:

- When Ambari Server is configured for HTTPS: `https://<your.ambari.server>:8443`
- When Ambari Server is configured for HTTP: `http://<your.ambari.server>:8080`

where <your.ambari.server> is the name of your ambari server host. For example, `c7401.ambari.apache.org`.



**Important:** Refresh your browser so that it loads the new version of the Ambari Web code. If you have problems, clear your browser cache manually, then restart Ambari Server.

18. Log in using the Ambari administrator credentials that you have set up. For example, the default name/password is `admin/admin`. You will see a Restart indicator next to each service after upgrading. Ambari upgrade has modified the configuration properties of your cluster based on the new configuration types and properties being made available for each service with this release of Ambari. Review these changes by comparing the previous configuration with the latest version created by "ambari-upgrade".



**Caution:** Do not manually restart these services unless future steps in the upgrade guide prompt you to do so. Manually restarting these services may significantly disrupt your upgrade. Ambari will restart each service automatically during the HDP upgrade.

19. If you have configured Ambari to authenticate against an external LDAP or Active Directory, you must re-run `ambari-server setup-ldap`.

20. If you have configured your cluster for Hive, Ranger or Oozie with an external database (Oracle, MySQL or PostgreSQL), you must run the following command to get the JDBC driver jar file in place: `ambari-server setup --jdbc-db` and `--jdbc-driver`.

A sample command to setup MySQL connector: `ambari-server setup --jdbc-db=mysql --jdbc-driver=/path/to/mysql/mysql-connector-java.jar_`

Ambari is ready and upgraded to 7.1.x.

## Upgrade to Ambari 7.1.0.0

After reviewing the information related to behavioral changes, Ambari UI and the Quick Links, and backing up the Ambari Server configuration file, perform the recommended steps for upgrading Ambari.

### Procedure

1. Turn off Auto Restart from Ambari web UI by browsing to Admin > Service Auto Start. Set Auto-Start Services to Disabled. Click Save.
2. If you are running SmartSense in your cluster, stop the service. From Ambari Web, browse to Services > SmartSense and select Stop from the Service Actions menu. Then, select Turn on Maintenance Mode from the Service Actions menu.

3. If you are running Ambari Metrics in your cluster, put it in Maintenance Mode. Don't stop the service yet. From Ambari Web, browse to Services > Ambari Metrics and select Turn on Maintenance Mode from the Service Actions menu.

a) Download the pre-upgrade script

1. SSH into the host running the Metrics Collector. Note, that in case Ambari Metrics System (AMS) is running in High Availability mode, there are multiple hosts running the Metrics Collector. These steps must be done on each Metrics Collector host.
2. Become the AMS service user. For example, run the following su command on Linux: `$ sudo su - ams`
3. Change to the /tmp directory.
4. Execute the following command to download the pre-upgrade script:

```
$ wget https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambaridc/centos7/7.x/updates/7.1.0.0/tars/ambari/ambari-metrics-pre-upgrade-7.1.0.0.138.tar.gz
```

To find all the download links in a centralised location, see [Software download matrix](#)

b) Extract and run the pre-upgrade script

1. Extract the pre-upgrade tarball: `$ tar xvzf ambari-metrics-pre-upgrade-<version>.tar.gz`
2. Change to the ambari-metrics-pre-upgrade-<version> directory. `$ cd ambari-metrics-pre-upgrade-<version>`
3. Make the pre-upgrade script executable: `$ chmod +x amms-pre-upgrade-2.6-to-7.x-<version>.sh`
4. Run the pre-upgrade script on all Metrics Collector hosts: `$ ./ams-pre-upgrade-2.6-to-7.x-<version>.sh`

Ambari Metrics Collector will shut down automatically.

c) Stop the service. From Ambari Web, browse to Services > Ambari Metrics and select Stop from the Service Actions menu.

4. If you are running Log Search in your cluster, stop the service. From Ambari Web, browse to Services > Log Search and select Stop from the Service Actions menu. Then, select Turn on Maintenance Mode from the Service Actions menu.
5. Stop the Ambari Server. On the host running Ambari Server: `ambari-server stop`
6. Stop all Ambari Agents. On each host in your cluster running an Ambari Agent: `ambari-agent stop`.
7. Back up your existing Ambari Server database. For example, to back up the default, embedded postgres db for Ambari:

```
mkdir /root/backups/
pg_dump -U ambari ambari > /root/backups/ambari-before-upgrade.sql
```

The default password is: bigdata.



**Note:** You must take a back up of the Ambari database if you have hosted the Ambari database on any other supported databases like MySQL or MariaDB.

8. Upgrade your Ambari Server database to a current supported version, only if your Ambari Server database version is not supported on the target Ambari version. To verify current supported versions, see the [Cloudera Support Matrix](#). For example, if you are running postgres v9.4, you must upgrade postgres to v9.6 or v10.2. If you are running MariaDB 10.2, upgrading the Ambari Server database is not required.

See the documentation for your database type and version for specific database upgrade instructions.

- Retrieve the new Ambari repo and replace the old repository file with the new repository file on all hosts in your cluster.



**Important:** Check your current directory before you download the new repository file to make sure that there are no previous versions of the `ambaridc.repo` file. If you do not, and a previous version exists, the new download is saved with a numeric extension, such as `ambaridc.repo.1`. Make sure that you copy the new version.

Select the repository appropriate for your environment from the following list:

- For RHEL or CentOS:

```
wget -nv https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/
p/ambaridc/centos7/7.x/updates/7.1.0.0/
ambaridc.repo -O /etc/yum.repos.d/ambaridc.repo
```

- For Ubuntu:

```
wget -nv https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com
/p/ambaridc/ubuntu18/7.x/updates/7.1.0.0/
ambaridc.list -O /etc/apt/sources.list.d/ambaridc.list
```

- Upgrade Ambari Server. On the host running Ambari Server:

For RHEL or CentOS:

```
yum clean all
```

```
yum info ambari-server
```

In the information output, visually validate that there is an available version containing "7.1.0.0"

```
yum upgrade ambari-server
```

For Ubuntu 18:

```
apt-get clean all
```

Before upgrading Ambari Server, you must update the username and password in the `ambari.list` file. Run the following command:

```
vi /etc/apt/sources.list.d/ambaridc.list
```

For example, the output displays the following:

```
#VERSION_NUMBER=7.1.0.0-138
#json.url = http://public-repo-1.hortonworks.com/HDP/hdp_urlinfo.json
deb https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/ambar
idc/ubuntu18/7.x/updates/7.1.0.0 Ambari main
apt-get update
apt-cache show ambari-server | grep Version
```

In the information output, visually validate that there is an available version containing "7.1.0.0"

```
apt-get install ambari-server
```

11. Check for upgrade success by noting progress during the Ambari Server installation process you started in Step 10.

- As the process runs, the console displays output similar to the following:  
Setting up Upgrade Process Resolving Dependencies Running transaction check
- If the upgrade fails, the console displays output similar to the following:  
Setting up Upgrade Process No Packages marked for Update
- A successful upgrade displays output similar to the following:  
Updated: ambari-server.noarch 0:7.1.0.0. Complete!

12. Upgrade all Ambari Agents. On each host in your cluster running an Ambari Agent:

- For RHEL or CentOS: Before upgrading Ambari Agent, you must update the username and password in the `ambaridc.repo` file. Run the following command:

```
vi /etc/yum.repos.d/ambaridc.repo
```

For example, the output displays the following:

```
#VERSION_NUMBER=7.1.0.0-138
[ambari-7.1.0.0]
#json.url = http://public-repo-1.hortonworks.com/HDP/hdp_urlinfo.json
name=ambari Version - ambari-7.1.0.0
baseurl=https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p
/ambaridc/centos7/7.x/updates/7.1.0.0
gpgcheck=1
gpgkey=https://[***USERNAME***]:[***PASSWORD***]@archive.cloudera.com/p/
ambaridc/centos7/7.x/updates/7.1.0.0/RPM-GPG-KEY/RPM-GPG-KEY-Jenkins
enabled=1
priority=1
```

```
yum upgrade ambari-agent
```

- For Ubuntu:

```
apt-get update
apt-get install ambari-agent
```

13. After the upgrade process completes, check each host to make sure the new files have been installed:

For CentOS 7.7:

```
rpm -qa | grep ambari-agent
```

For Ubuntu 18:

```
dpkg -l ambari-agent
```

14. Upgrade Ambari Server database schema. On the host running Ambari Server: `ambari-server upgrade`. When the Ambari Server database schema has been upgraded, you should see command output like this: Ambari Server 'u pgrade' completed successfully.



**Note:** If the Ambari-server upgrade fails with the error: NO JSON Object could be decoded, then you must check `pg_hba.conf`.



15. Start the Ambari server from the host running using the following command:

```
ambari-server start
```



**Note:** If Ambari-server fails to start due to 'FATAL: role "mapred" does not exist', then you must note that if an embedded postgres was upgraded, ambari-server setup should be run again.

16. Start all Ambari Agents on each host in your cluster running an Ambari Agent using the following command:

```
ambari-agent start
```

17. Open Ambari Web UI. Point your browser to the Ambari Web UI:

- When Ambari Server is configured for HTTPS: `https://<your.ambari.server>:8443`
- When Ambari Server is configured for HTTP: `http://<your.ambari.server>:8080`

where <your.ambari.server> is the name of your ambari server host. For example, `c7401.ambari.apache.org`.



**Important:** Refresh your browser so that it loads the new version of the Ambari Web code. If you have problems, clear your browser cache manually, then restart Ambari Server.

18. Log in using the Ambari administrator credentials that you have set up. For example, the default name/password is admin/admin. You will see a Restart indicator next to each service after upgrading. Ambari upgrade has modified the configuration properties of your cluster based on the new configuration types and properties being made available for each service with this release of Ambari. Review these changes by comparing the previous configuration with the latest version created by "ambari-upgrade".



**Caution:** Do not manually restart these services unless future steps in the upgrade guide prompt you to do so. Manually restarting these services may significantly disrupt your upgrade. Ambari will restart each service automatically during the HDP upgrade.

19. If you have configured Ambari to authenticate against an external LDAP or Active Directory, you must re-run `ambari-server setup-ldap`.

20. If you have configured your cluster for Hive, Ranger or Oozie with an external database (Oracle, MySQL or PostgreSQL), you must run the following command to get the JDBC driver jar file in place: `ambari-server setup --jdbc-db` and `--jdbc-driver`.

A sample command to setup MySQL connector: `ambari-server setup --jdbc-db=mysql --jdbc-driver=/path/to/mysql/mysql-connector-java.jar_`

Ambari is ready and upgraded to 7.1.x.

## Download cluster blueprints

Ensure you download the cluster blueprints before you begin the upgrade.

### About this task

#### Procedure

1. To download the cluster blueprints with hosts, do the following:

```
curl ${securecurl} -H "X-Requested-By: ambari" -X GET -u ${ambariuser}:${ambaripwd} ${ambariprotocol}://${ambariserver}:${ambariport}/api/v1/clusters/${clustername}?format=blueprint_with_hosts > "${backupdir}/${clustername}_blueprint_with_hosts_"$(date +%Y%m%d%H%M%S)".json
```

2. To download the cluster blueprints without hosts, do the following:

```
curl ${securecurl} -H "X-Requested-By: ambari" -X GET -u ${ambariuser}:${ambaripwd} ${ambariprotocol}://${ambariserver}:${ambariport}/api/v1/clu
```

```
sters/${clustername}?format=blueprint > "${backupdir}/${clustername}_bl
ueprint_without_hosts_"$(date +%Y%m%d%H%M%S)".json
```

where,

securecurl = -k for https

securecurl = -k for https

ambariuser = Ambari User Name

ambaripwd = Ambari Password

ambariprotocol = http or https

ambariserver = Ambari Server Name

ambariport = Ambari Port

clustername = The Name of the cluster

backupdir = The folder to download the blueprint

## Mandatory Post-Upgrade Tasks

You must ensure to complete the post-upgrade tasks after upgrading to Ambari 7.1.x.x.

Depending on the configuration of your cluster and your current Ambari version, you must upgrade any of the following features in your cluster, as described in the following topics:

### Upgrading Ambari Metrics System and SmartSense

As part of the Ambari 7.1.x.x release, the schema for the Ambari Metrics System has significantly changed. This change requires additional steps to be performed before upgrading to HDP intermediate bits. Because SmartSense depends on the Ambari Metrics System, SmartSense also requires similar handling. The following steps guide you through upgrading the Ambari Metrics System and SmartSense.

The following steps will guide you through upgrading the Ambari Metrics System and SmartSense.



**Note:** Before following these steps, ensure that the Ambari Metrics System and SmartSense are stopped and in Maintenance mode. The HDP intermediate bits version should be registered and the packages should be installed.

#### Upgrading Ambari Metrics

You must stop Ambari Metrics and then run commands to upgrade Metrics Monitor, Metrics Collector, and Grafana on the cluster nodes where the relevant components are running.

#### About this task

#### Procedure

1. Confirm that the Ambari Metrics service is stopped and in Maintenance Mode. If Ambari Metrics service is not stopped, from Ambari Web browse to Services > Ambari Metrics and select Stop from the Service Actions menu.

2. On every host in your cluster that is running a Metrics Monitor, run the following commands:

For RHEL/CentOS:

- `yum clean all`
- `yum upgrade ambari-metrics-monitor`
- `yum upgrade ambari-metrics-hadoop-sink`

For Ubuntu 18:

- `apt-get clean all`
- `apt-get update`
- `apt-get install ambari-metrics-assembly`

3. Execute the following command on all hosts running the Metrics Collector:

For RHEL/CentOS:

- `yum upgrade ambari-metrics-collector`

For Ubuntu 18:

- `apt-get clean all`
- `apt-get update`
- `apt-get install ambari-metrics-collector`

4. Execute the following command on the host running the Grafana component:

For RHEL/CentOS:

- `yum upgrade ambari-metrics-grafana`

For Ubuntu 18:

- `apt-get clean all`
- `apt-get update`
- `apt-get install ambari-metrics-grafana`

**Note:**

- Do not start the Ambari Metrics System service. Do not restart other services. The HDP upgrade process will restart all the services.
- The updated Ambari Metrics Sink jars will be installed on all hosts. During the upgrade process, after the restart, each service start uses the latest sink implementations.

## Backup Ambari-Metrics

### Procedure

1. Stop Ambari-Metrics service in Ambari UI.
2. Back-up services data to restore later. For Ambari-Metrics on the node where the `ambari-metrics-collector` is installed.

```
tar czf ams-backup.tar.gz /var/lib/ambari-metrics-collector/hbase
```



**Note:** Ensure that you check the AMS configuration and the HBase root directory to determine the backup directory.

## Upgrading SmartSense

Make sure you stop SmartSense and Ambari Metrics. Then you must upgrade the HST binaries, upgrade Ambari service and Ambari view, upgrade the HST Gateway if installed, and then restart the Ambari server.

### About this task

## Procedure

1. Confirm that SmartSense and Ambari Metrics are stopped and in Maintenance Mode.
  - If SmartSense is not stopped, from Ambari Web browse to Services > SmartSense and select Stop from the Service Actions menu.
  - If Ambari Metrics is not in Maintenance Mode, From Ambari Web browse to Services > Ambari Metrics and select Turn On Maintenance Mode from the Service Actions menu.
2. Upgrade binaries on the HST server and all HST agents on every node in the cluster, assuming that the Ambari repository is configured on all nodes in the cluster:

RHEL or CentOS:

- `yum clean all`
- `yum info smartsense-hst`
- In the info output, visually validate that there is an available version containing "1.5.x": `yum upgrade smartsense-hst`

Ubuntu 18:

- `apt-get clean all`
- `apt-get update`
- `apt-cache show smartsense-hst | grep Version`
- In the info output, visually validate that there is an available version containing "1.5.x": `apt-get install smartsense-hst`

3. Upgrade Ambari service and Ambari view by running the `hst upgrade-ambari-service` command as the root user from the machine running the Ambari server. You can run the command in the interactive or non-interactive mode:

Interactive mode example:

```
# hst upgrade-ambari-service
Please enter Ambari Server hostname (ambari-server.
hortonworks.local):
Please enter Ambari Server port (8080):
Please enter Ambari admin user id (admin):
Please enter password for admin:

Un-installing old view ...
Installing new view ...
Removing deprecated alerts ...
Updating SmartSense configurations in Ambari ...

SmartSense service upgrade completed!
```



**Note:** It is required to restart Ambari Server for changes to reflect. Please restart ambari using `ambari-server restart`

Non-interactive mode example:

```
# hst upgrade-ambari-service -u admin -p 8080 -H ambari-server.hortonwor
ks.local -P MySecurePassword123
Un-installing old view ...
Installing new view ...
Removing deprecated alerts ...
Updating SmartSense configurations in Ambari ...
SmartSense service upgrade completed!
```



**Note:** It is required to restart Ambari Server for changes to reflect. Please restart ambari using `ambari-server restart`

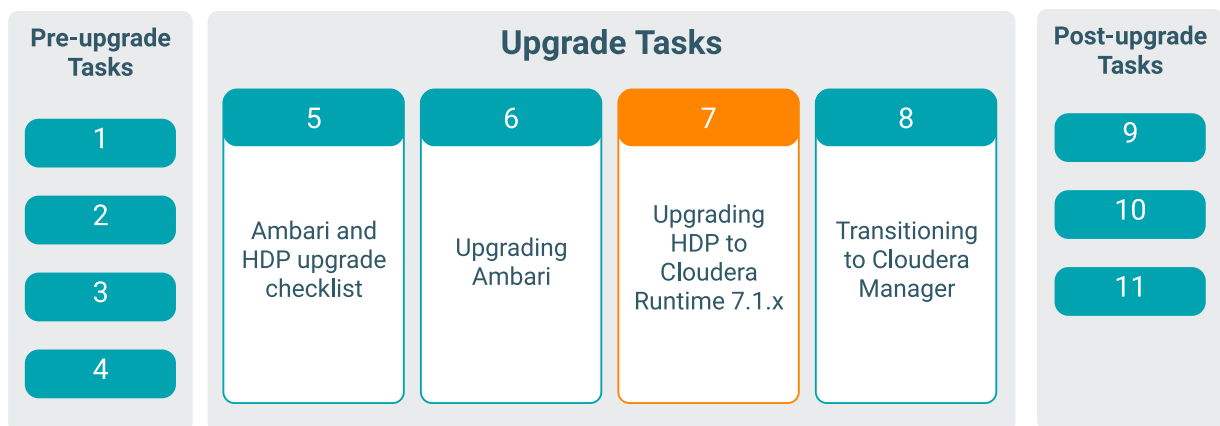
4. Restart the Ambari server: `# ambari-server restart`
5. If you have HST Gateway installed, you need to also upgrade your HST Gateway:
  - If the HST Gateway is installed on the same node as HST Server or HST Agent, then the HST Gateway will get upgraded along with them.
  - If the HST Gateway is a standalone node outside of the cluster, perform upgrade steps described in [Upgrading SmartSense Gateway](#).



**Important:** It is very important to make sure you do not start the SmartSense service. It will start automatically during the HDP upgrade process.

## Upgrading HDP to Cloudera Runtime 7.1.x

To complete the first stage of upgrading from HDP 2.6.5.x to CDP Private Cloud Base, you must upgrade to Cloudera Runtime 7.1.x.



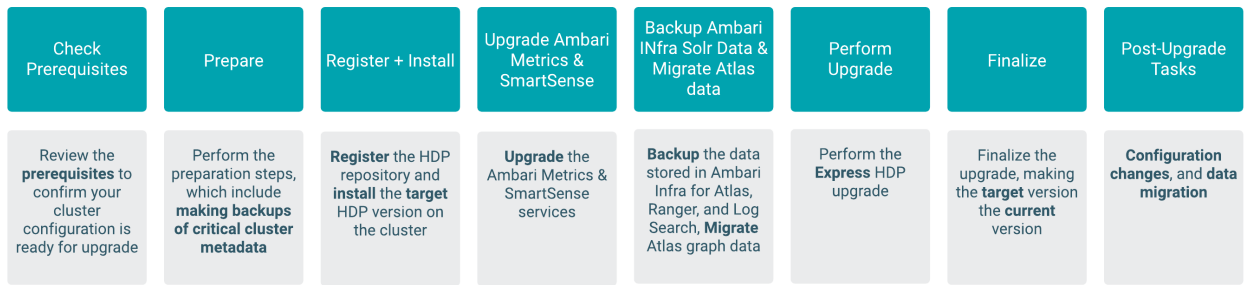
### HDP Prerequisites

Before upgrading from HDP 2.6.5.x to the Cloudera Runtime 7.1.x, you must review and complete the prerequisites.

### Upgrade process

Read through the complete information about the upgrade options, prerequisites, and the overall process before starting the upgrade. Cloudera recommends that you validate these steps in a test environment to adjust and account for any special configurations for your cluster.

The high-level process for performing an HDP intermediate bits upgrade is as follows:



Ambari will guide you through the steps required to upgrade the HDP intermediate bits. Make sure Ambari and the cluster are healthy, operating normally, and all service checks are passing.



**Note:** Be sure to review the available HDP intermediate bits upgrade scenarios below. Cloudera recommends that you first upgrade to Ambari 7.1.x.x before upgrading to the HDP intermediate bits unless otherwise noted. After upgrading Ambari, ensure that the cluster is operating normally and service checks are passing prior to attempting an HDP upgrade.

Ambari 7.1.x.x supports only the express method for upgrading HDP 2.6.5.x to HDP intermediate bits. An Express Upgrade orchestrates the HDP upgrade in an order that will incur cluster downtime.



**Attention:**

- Suspend all the Oozie workflows and cron jobs before upgrading.
- Stop all the applications that access the HDP Services.
- Anything that interfaces your cluster or access data must be completely stopped on your HDP 2.6.x cluster.

## Kerberos cluster

To perform an HDP intermediate bits upgrade using Ambari, your kerberized cluster must meet certain prerequisites.



**Important:** For a Kerberized cluster, you must review and configure your KDC:

- KDC Admin Credentials:

The Ambari Server adds new components as a part of the HDP 2.6.5.x to HDP intermediate bits upgrade and needs to be configured to enable password encryption and saving the KDC admin credentials so that necessary principals can be created. For steps on how to do this, see [Encrypt Database and LDAP Passwords in Ambari](#), [Managing Admin Credentials](#), and [Authentication and Kerberos Issues](#).

- KDC Admin Host FQDN:

Ensure that the KDC Admin host is set to a fully qualified domain name (FQDN) and not an IP address. You can check this setting by going to Services > Kerberos > Configs > KDC Hosts

- KDC kadmin principal:

The MIT KDC integration in Ambari 7.1.x.x has been improved to connect more securely with the Kerberos Administration server (kadmind). This increased security expects a Kerberos admin service principal to be present with a specifically formatted principal name. The format expected is kadmin/fully.qualified.kdc.hostname@REALM. This expectation is a behavior change from previous versions of Ambari, where having such a Kerberos admin service principal was not required. This principal is present by default in most MIT KDC installations, but some customers have reported that this principal does not exist in their KDC. Due to this, it is recommended that before upgrading a kerberized cluster to Ambari 7.1.x.x, you ensure this principal exists in your KDC database.

## Before upgrading any cluster

To perform an HDP intermediate bits upgrade using Ambari, your cluster must meet certain prerequisites. Meeting these prerequisites is essential for Ambari to know that the cluster is in a healthy operating mode and can successfully manage the upgrade process.

For any Cluster

**Disk Space:** Be sure to have adequate space on /usr/hdp for the target HDP version. Each complete install of an HDP version occupies about 5 GB of disk space.

**Ambari Agent Heartbeats:** All Ambari Agents must be communicating and heartbeating to Ambari Server. Any hosts that are not heartbeating must be in Maintenance mode.

**Hive Upgrade:** The upgrade process does not back up the Hive Metastore, nor does it compact ACID tables. Before upgrading Hive, you must:

- Manually back up your Hive metastore database after using the pre-upgrade tool, described later, and before upgrading.
- If you have ACID tables in your Hive metastore database, enable ACID operations using Ambari Web or set Hive configuration properties to enable ACID.



**Note:** You must remove the components that are unsupported in CDP Private Cloud Base. For more information, see [HDP Core component version changes](#). If you are not using a component in the CDP cluster, then you must remove them. The unused component extends the time and complexity of the HDP to CDP upgrade.

**Host Maintenance Mode**

The following two scenarios are checked:

- Any hosts in Maintenance mode must not be hosting any Service Master Components.
- Hosts running Worker Components display a warning if they are in Maintenance mode. You can proceed with the upgrade. But these hosts will not be upgraded. Before finalising the upgrade, you must delete these hosts from the cluster.

**Service Maintenance Mode:** No services can be in Maintenance mode, except for Ambari Metrics System, SmartSense, and Log Search.

**Services Started:** All services must be started, except for Ambari Metrics System, SmartSense, and Log Search.

**Service Checks:** All service checks must pass. Be sure to run Service Actions > Run Service Check on all services (and remediate if necessary) prior to attempting an HDP upgrade.

**All Prior Upgrades are Finalized:** Any prior upgrade started with Ambari must be finalized.

## Backup HDP Cluster

Cloudera recommends that you back up the HDP 2.6.5.x services managed by Ambari 2.6.x.x.



**Note:** These backups are required and helps you incase in the middle of the upgrade to HDP 7.1.x you decide to downgrade the cluster back to original state.

## Backup and Restore Databases

The steps for backing up and restoring databases differ depending on the database vendor and version that you select for your cluster. It is beyond the scope of this document.



**Important:** Restore the databases to their exact state as of when you took the backup. Do not merge any changes that may have occurred during the subsequent upgrade.

See the following vendor resources for more information:

- For MariaDB 5.5, see [MariaDB](#)
- For MySQL 5.5, see [MySQL5.5](#)
- For MySQL 5.6, see [MySQL 5.6](#)
- For MySQL 5.7, see [MySQL 5.7](#)
- For PostgreSQL 8.4, see [PostgreSQL 8.4](#)

- For PostgreSQL 9.3, see [PostgreSQL 9.3](#)
- For Oracle 11gR2, see [Oracle 11gR2](#)

## Backup Ranger

You must stop the Ranger Admin and Ranger KMS Service and its dependent services.

### Backup Ranger Admin Database

You must execute the below commands on the Ranger database host.

MySQL

1. First create directory (mkdir -p /root/backups/ranger/db/admin)

```
$ mysqldump -u [username] -p db_name > dump_fileName.sql
```

or specify destination as

```
/some_dir/dump_fileName.sql
```

2. Select the Enter key.
3. Type the database password at the password prompt.

Example

```
mysqldump -u rangeradmin -p ranger > /root/backups/ranger/db/admin/ranger.sql
```

POSTGRES

1. Dump the contents of a database to a file by running the following command. Replace database name with the name of the database to be backed up:

```
pg_dump -U username dbname > dbname.bak
```

Example

```
pg_dump -U rangeradmin ranger > /root/backups/ranger/db/admin/db/ranger.sql
```

2. Type the database password at the password prompt.

### Backup Ranger KMS Database

You must backup the MySQL, POSTGRES, and Oracle databases.



**Note:** It is a good practice to note down the master key password before backing up the databases.

MySQL

1. Create directory (mkdir /etc/ranger/hdp2\_backup/db/kms)
2. \$ mysqldump -u [username] -p db\_name > dump\_fileName.sql (or specify destination as /some\_dir/dump\_fileName.sql)
3. Select the Enter key.
4. Type the database password at the password prompt.

Example

```
mysqldump -u rangerkms -p rangerkms > /root/backups/ranger/db/kms/rangerkms.sql
```

POSTGRES



1. Dump the contents of a database to a file by running the following command. Replace dbname with the name of the database to be backed up:

```
pg_dump -U [username for database] rangerkms > rangerkms.sql
```

Example

```
pg_dump -U rangerkms rangerkms > /root/backups/ranger/db/kms/rangerkms.sql
```

Oracle

1. Set path to Oracle home :
  - export ORACLE\_HOME=/opt/oracle/product/12.2.0
  - export PATH=\${PATH}:\${ORACLE\_HOME}/bin
  - export ORACLE\_SID=orcl12c (db\_name)

2. Backup Ranger admin database:

```
exp userid=rangeradmin/rangeradmin owner=rangeradmin log=backups/ranger/db/admin/admin_db_bkp.log file=backups/ranger/db/admin/orcl12c.sql statistics=none
```

3. Backup Ranger KMS database:

```
exp userid=rangerkms/rangerkms owner=rangerkms log=backups/ranger/db/kms/kms_db_bkp.log file=backups/ranger/db/kms/orcl12c.sql statistics=none
```

## Backup Atlas

You can back up Atlas data by backing-up HBase tables.

### Procedure

1. Stop Atlas service from Ambari.
2. Setup Kerberos credentials. Locate the atlas user's keytab and use kinit to cache the kerberos ticket for atlas user. For example,

```
kinit -kt path/to/atlas.service.keytab atlas/hostname@domain
```

### Backup HBase tables

Follow these steps to back up atlas\_titan table.

### Procedure

1. Go to HBase Shell.
2. Disable atlas\_janus and ATLAS\_ENTITY\_AUDIT\_EVENTS tables.
3. Create a snapshot of the atlas\_janus and ATLAS\_ENTITY\_AUDIT\_EVENTS HBase tables.
4. (Optional step) If the backup needs to be preserved on the local file system, then use the additional command to copy it locally.

```
hbase org.apache.hadoop.hbase.snapshot.ExportSnapshot -snapshot 'atlas-titan-backup' -copy-to /path/to/backup/location
```

The above steps can be performed using the following commands from HBase shell.

```
> disable 'atlas_titan'
> snapshot 'atlas_titan', 'atlas-titan-backup'
```

```
> enable 'atlas_titan'
> exit
```

```
> disable 'ATLAS_ENTITY_AUDIT_EVENTS'
> snapshot 'ATLAS_ENTITY_AUDIT_EVENTS', 'ATLAS_ENTITY_AUDIT_EVENTS-backup'
> enable 'ATLAS_ENTITY_AUDIT_EVENTS'
> exit
```

## Backup Ambari Infra Solr

Back up all your ambari infra solr collections. Run the Solr APIs to backup the collections individually. You must create the backup directory on each Ambari Infra Solr node. However, the backup command takes the backup of shards only on the nodes where they reside.

### About this task

The specified `/path/to/backup/` directory must be created first on the host where the ambari-infra-solr is installed. The specified location needs to be accessible to the user by executing the command.

```
mkdir -p /path/to/backup/directory
chown -R infra-solr:hadoop /path/to/backup
```

For unsecured cluster

```
#Atlas infra-solr collections
curl -v
"http://$INFRA_SOLR_URL/solr/vertex_index/replication?command=BACKUP&name=vertex_index_backup&location=/path/to/backup/directory"

curl -v
"http://$INFRA_SOLR_URL/solr/edge_index/replication?command=BACKUP&name=edge_index_backup&location=/path/to/backup/directory"

curl -v
"http://$INFRA_SOLR_URL/solr/fulltext_index/replication?command=BACKUP&name=fulltext_index_backup&location=/path/to/backup/directory"

#Ranger audit infra-solr collections
curl -v
"http://$INFRA_SOLR_URL/solr/ranger_audits/replication?command=BACKUP&name=ranger_audits_backup&location=/path/to/backup/directory"
#Log Search infra-solr collections
curl -v
"http://$INFRA_SOLR_URL/solr/hadoop_logs/replication?command=BACKUP&name=hadoop_logs_backup&location=/path/to/backup/directory"

curl -v
"http://$INFRA_SOLR_URL/solr/audit_logs/replication?command=BACKUP&name=audit_logs_backup&location=/path/to/backup/directory"

curl -v
"http://$INFRA_SOLR_URL/solr/history/replication?command=BACKUP&name=history_backup&location=/path/to/backup/directory"
```

For secured cluster

If the cluster is Kerberized, then you must kinit as the service principal.

```
#Atlas infra-solr collections
curl -v --negotiate -u:
"http://$INFRA_SOLR_URL/solr/vertex_index/replication?command=BACKUP&name=vertex_index_backup&location=/path/to/backup/directory"
```

```

curl -v --negotiate -u:
"http://$INFRA_SOLR_URL/solr/edge_index/replication?command=BACKUP&name=ed
ge_index_backup&location=/path/to/backup/directory"

curl -v --negotiate -u:
"http://$INFRA_SOLR_URL/solr/fulltext_index/replication?command=BACKUP&na
me=fulltext_index_backup&location=/path/to/backup/directory"

#Ranger audit infra-solr collections
curl -v --negotiate -u:
"http://$INFRA_SOLR_URL/solr/ranger_audits/replication?command=BACKUP&name=r
anger_audits_backup&location=/path/to/backup/directory"
#Log Search infra-solr collections
curl -v --negotiate -u:
"http://$INFRA_SOLR_URL/solr/hadoop_logs/replication?command=BACKUP&name=ha
dooop_logs_backup&location=/path/to/backup/directory"

curl -v --negotiate -u:
"http://$INFRA_SOLR_URL/solr/audit_logs/replication?command=BACKUP&name=a
udit_logs_backup&location=/path/to/backup/directory"

curl -v --negotiate -u:
"http://$INFRA_SOLR_URL/solr/history/replication?command=BACKUP&name=hist
ory_backup&location=/path/to/backup/directory"

```

It is highly recommended to save the data folder for disaster recovery.

### Procedure

1. Stop Ambari-Infra service in Ambari UI.
2. Back-up services data to restore later. For Ambari-Infra on the node where the ambari-infra-solr is installed.

```
tar -czf ambari-infra-backup.tar.gz /var/lib/ambari-infra-solr/data
```

### Backup Hive

To backup the Hive service, you must backup the Hive database and create a snapshot of the Hive warehouse.

Backup Hive table data using a snapshot

In Ambari, go to Services > Hive > Configs, and check the value of hive.metastore.warehouse.dir to determine the location of the Hive warehouse, /apps/hive/warehouse by default.

1. On any node in the cluster, as the HDFS superuser, enable snapshots. For example,

```
$ sudo su - hdfs
$ hdfs dfsadmin -allowSnapshot /apps/hive/warehouse
```

Output is:

```
Allowing snapshot on /apps/hive/warehouse succeeded
```

2. Create a snapshot of the Hive warehouse. For example,

```
$ hdfs dfs -createSnapshot /apps/hive/warehouse
```

Output includes the name and location of the snapshot:

```
Created snapshot /apps/hive/warehouse/.snapshot/s20210308-104702.892
```

Back up Hive Metastore

On the node where the database you use for Hive Metastore resides, back up Hive Metastore before upgrading to HDP. For example, in MySQL, dump each database as follows:

```
mysqldump -u <hive_db_user> <hive_db_schema_name> > </path/to/dump_file>
```

If you use another database for the Hive Metastore, use the equivalent command, such as `pg_dump` for Postgres to dump the database.

## Backup HBase

The rollback procedure rolls back HDFS and, the data in HBase is also rolled back automatically as part of this procedure. In addition, HBase metadata stored in ZooKeeper is recovered as part of the ZooKeeper rollback procedure.

### About this task

If your cluster is configured to use HBase replication, Cloudera recommends that you document all replication peers. If necessary (for example, because the HBase znode has been deleted), you can roll back HBase as part of the HDFS rollback without the ZooKeeper metadata. This metadata can be reconstructed in a fresh ZooKeeper installation, with the exception of the replication peers, which you must add back.

### Procedure

Stop HBase before backing up HDFS.

## Backup Kafka

Before you begin the upgrade process, you need to explicitly set the Kafka protocol version to match what's being used currently among the brokers and clients. Update `server.properties` on all brokers as follows:

### Procedure

1. Log in to Ambari.
2. Choose the Kafka service.
3. Select the Configs page.
4. Find the Custom kafka-broker section.
5. Add the following properties:
  - `inter.broker.protocol.version=current_Kafka_version`
  - `log.message.format.version=current_Kafka_version`

Replace `current_Kafka_version` with the version of Apache Kafka currently being used, in case of HDP-2.6.5 it is 1.0.

6. Save your changes. The information is automatically copied to each broker.



**Important:** Once the Kafka log format and protocol version configurations (the `inter.broker.protocol.version` and `log.message.format.version` properties) are set to the new version (or left blank, which means to use the latest version), Kafka rollback is not possible.

## Backup Oozie

To backup the Oozie service, you must backup the Oozie database.

For other databases, Cloudera recommends that you perform backup of your databases before beginning the upgrade.

- Ambari database
- Hive Metastore database (after the pre-upgrade tool runs compaction)
- Oozie Server database
- Ranger database

- Ranger KMS database
- Druid database
- Superset Database

### Backup Knox

With the backup and rollback of Ambari, Knox is also backed up and rolled back by default.

### Backup Logsearch

With the backup of Ambari-Infra, Logsearch is also backed up by default.

### Backup Zeppelin

With the backup and rollback of HDFS, Zeppelin is also backed up and rolled back by default.

### Backup HDFS

You can roll back an upgrade from CDP Private Cloud Base 7 to HDP 2. The rollback restores your HDP cluster to the state it was in before the upgrade. This means any change that happened after taking the backups as instructed from points 4 to 7 will be reverted from the HDFS cluster. You must perform additional steps before upgrading your HDP cluster if you have configured the NameNode HA.

#### Procedure

1. If you have configured the NameNode HA, then locate the Active NameNode from Ambari Web > Services > HDFS in the Summary area.
2. Check the NameNode directory to ensure that there is no snapshot of any prior HDFS upgrade. Specifically, using Ambari Web, browse to Services > HDFS > Configs, and examine the `dfs.namenode.name.dir` in the NameNode Directories property. Make sure that only a `/current` directory exists and remove `/previous` directory on the NameNode host.
3. Create the following log and other files. The commands below back up additional state from the file system in addition to what `fsImage` contains. This information can be used to validate the filesystem state after the upgrade

by comparing the output before or after the upgrade. This is helpful, but it is not required for troubleshooting any issues that arise during the upgrade.



**Note:** The `fsck` or `ls` command may take a significant amount of time if the number of blocks or files in HDFS is in the ten to a hundred million or more range.

Authenticate as the HDFS user `su -l [HDFS_USER]`



**Note:** If the cluster is Kerberized, then you must kinit as the service principal. If the cluster is not Kerberized, then you must run as a cluster administrator user or log in as the service user.

Run the following, where `[HDFS_USER]` is the HDFS Service user. For example, `hdfs`:

- a) Run `fsck` with the following flags and send the results to a log. The resulting file contains a complete block map of the file system. Use this log later to confirm the upgrade.

```
hdfs fsck / -files -blocks -locations > dfs-old-fsck-1.log
```

- b) Create a list of all the DataNodes in the cluster.

```
hdfs dfsadmin -report > dfs-old-report-1.log
```

- c) Capture the complete namespace of the file system. The following command does a recursive listing of the root file system:

```
hdfs dfs -ls -R / > dfs-old-lsr-1.log
```

- d) Take a backup of the HDFS data to the backup instance of your HDFS, if you have such a system.

4. Create a backup from the configuration directory under `/etc/hadoop/conf` into a backup directory on all of your hosts.

5. Save the namespace. As the HDFS user, "`su -l [HDFS_USER]`", you must put the cluster in Safe Mode.

```
hdfs dfsadmin -safemode enter
hdfs dfsadmin -saveNamespace
```



**Note:** In a highly-available NameNode configuration, the command `hdfs dfsadmin -saveNamespace` sets a checkpoint in the first NameNode specified in the configuration, in `dfs.ha.namenodes.[nameserviceID]`. You can also use the `dfsadmin -fs`.

6. Create a backup from the directory located in `$(dfs.namenode.name.dir)/current` into a backup directory.
7. Create a backup from the directory located in `$(dfs.journalnode.edits.dir)/current` into a backup directory.
8. As the HDFS user, `su -l [HDFS_USER]`, take the NameNode out of Safe Mode. `hdfs dfsadmin -safemode leave`  
From this point, you have all the data that might be required to restore HDFS to the state where it was at the time when HDFS entered safe mode.

A few points to consider regarding the backups:

- `fsck ls`, and other command outputs happened before entering safe mode if you went step by step over the guide. The guide initiates safe mode before the commands that require safe mode, however you can enter safe mode earlier.
- If you come out of safe mode, and leave the cluster running, then you are allowing mutations on the filesystem that are not saved in the backup.

## Backup ZooKeeper

You must backup Zookeeper.

### Procedure

1. Stop Zookeeper service in Ambari UI.

2. Back-up services data to restore later. For Zookeeper on every node.

```
tar -czf backup.tar.gz /hadoop/zookeeper
```

## Backup Ambari

You must backup Ambari.

### Procedure

1. Create Ambari database dump.

- a) On Postgres

```
pg_dump ambari -U ambari > pre_upgrade.sql
```

- b) On MySQL or MariaDB

```
mysqldump -u ambari ambari > pre_upgrade.sql
```

2. Backup the following:

```
/etc/ambari-server/conf/ambari.properties  
/etc/ambari-agent/conf/ambari-agent.ini
```

## Backup databases

To continue upgrading to Cloudera Runtime 7.1.x, you must have completed the HDP prerequisites tasks. Ensure you backup your databases before beginning the upgrade.

### Before you upgrade

Make sure that you have reviewed and completed all prerequisites described in above.

Cloudera recommends that you perform backup of your databases before beginning the upgrade.

- Ambari database
- Hive Metastore database (after the pre-upgrade tool runs compaction)
- Oozie Server database
- Ranger database
- Ranger KMS database
- Druid database
- Superset Database

### Checkpoint HDFS

You must perform additional steps before upgrading your HDP cluster if you have configured NameNode HA to create a checkpoint for HDFS.

### Procedure

1. If you are configured for NameNode HA, then locate the Active NameNode from Ambari Web > Services > HDFS in the Summary area.
2. Check the NameNode directory to ensure that there is no snapshot of any prior HDFS upgrade. Specifically, using Ambari Web, browse to Services > HDFS > Configs, and examine the `dfs.namenode.name.dir` in the NameNode Directories property. Make sure that only a `/current` directory exists and remove `/previous` directory on the NameNode host.
3. Create the following log and other files. The commands below back up additional state from the file system in addition to what `fsImage` contains. This information can be used to validate the filesystem state after the upgrade

by comparing the output before or after the upgrade. This can help but is not required for troubleshooting if issues arise during the upgrade.



**Note:** The `fsck` or `ls` command may take a significant amount of time if the number of blocks or files in HDFS is in ten to a hundred million range.

Authenticate as the HDFS user `su -l [HDFS_USER]`



**Note:** If the cluster is Kerberized, then you must kinit as the service principal. If the cluster is not Kerberized, then you must run as a cluster administrator user or log in as the service user.

Run the following (where `[HDFS_USER]` is the HDFS Service user, for example, `hdfs`):

- Run `fsck` with the following flags and send the results to a log. The resulting file contains a complete block map of the file system. You use this log later to confirm the upgrade.

```
hdfs fsck / -files -blocks -locations > dfs-old-fsck-1.log
```

- Create a list of all the DataNodes in the cluster.

```
hdfs dfsadmin -report > dfs-old-report-1.log
```

- Capture the complete namespace of the file system. The following command does a recursive listing of the root file system:

```
hdfs dfs -ls -R / > dfs-old-lsr-1.log
```

- Take a backup of the HDFS data to your local file system or the backup instance of your HDFS.

#### 4. Save the namespace. As the HDFS user, "`su -l [HDFS_USER]`", you must put the cluster in Safe Mode.

- `hdfs dfsadmin -safemode enter`
- `hdfs dfsadmin -saveNamespace`



**Note:** In a highly-available NameNode configuration, the command `hdfs dfsadmin -saveNamespace` sets a checkpoint in the first NameNode specified in the configuration, in `dfs.ha.namenodes.[nameserviceID]`. You can also use the `dfsadmin -fs`.

#### 5. Copy the checkpoint files located in `$(dfs.namenode.name.dir)/current` into a backup directory.

#### 6. Store the layout version for the NameNode located at `$(dfs.namenode.name.dir)/current/VERSION` into a backup directory where `$(dfs.namenode.name.dir)` is the value of the config parameter `NameNode directories`.

#### 7. This file will be used later to verify that the layout version is upgraded. As the HDFS user, `su -l [HDFS_USER]`, take the NameNode out of Safe Mode. `hdfs dfsadmin -safemode leave`

## Register software repositories

Before you use Ambari to perform the stack upgrade, you must register the software repositories for the new target version with Ambari and then install the software on all hosts in the cluster.

### Procedure

1. Login to Ambari.
2. Browse to `Cluster Admin Stack and Versions`.
3. Click the `Versions` tab. You see the version currently running, marked as `Current`.



**Note:** The full version depends on the HDP version you are actually running. For example, if you are currently running the HDP 2.6.5.x release, you would see something like `HDP-2.6.5.0-292` as the full version number.

4. Click `Manage Versions`.
5. Proceed to register a new version by clicking `Register Version`.



6. Select the software version and method of delivery for your cluster.

a) Choose HDP Stack

Available HDP minor versions display on tabs. When you click a tab, Ambari displays available maintenance versions for that HDP stack in a drop-down list. When you click a specific maintenance version, a list of available services and their version is displayed in a table.

b) Choose the HDP Version

If your Ambari host has internet access, available maintenance versions display as options in a drop-down list.

Make sure to use the Version Definition File (VDF) option for HDP 7.1.x.x.

If you see the VDF option for another version, click Add Version... and upload the VDF file for HDP 7.1.x.x.

If you have difficulties, contact Cloudera Support to obtain the VDF for HDP 7.1.x.x.

c) Choose Repository Delivery Method

- Using a public repository requires internet connectivity. To use the public software repositories, see the list of HDP intermediate bits repositories for each operating system.
- HDP intermediate bits Repositories are available for HDPDC, HDPDC-GPL, and HDP-UTILS. To use the repositories, see [HDP intermediate bits Repositories for 7.1.x.0 version](#).
- Using a local repository requires that you have already configured the software in a repository available in your network. If you are using a local repository, enter the Base URLs for your local repository.

To find all the download links in a centralised location, see [Software Download Matrix](#).

d) If you are on SLES12, edit the repository URLs:

- For each URL, remove the basicAuth credentials between `https://` and `archive.cloudera.com`
- Add `?credentials=ambari.cat` postfix at the end of each URL.
- For example, the base URL must be this: `https://archive.cloudera.com/p/HDPDC/sles12/7.x/updates/7.1.4.0?credentials=ambari.cat`

e) Select the Skip Repository Base URL validation (Advanced) checkbox.

7. Click Save.

8. Click Dashboard.



**Note:** For SLES 12:

- Create the `hdp.cat` file with your paywall credentials. The file is available here: `/etc/zypp/credentials.d/hdp.cat`. This is the default location. If there is a change in the location of the `hdp.cat` file, then you must provide the changed location. You must update the username and password in the `hdp.cat` file:

```
username=<username>
```

```
password=<password>
```

- Add `?credentials=hdp.cat` as a postfix to the end of BaseURL and `gpgkey`. For example, `https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/ambari.repo?credentials=hdp.cat`
- For more information, see [Zypp Configuration options](#).

## HDP Intermediate bits for 7.1.x.0 Repositories

HDP repositories are available for RHEL/CentOS, Sles, and Ubuntu.



**Note:** For repositories that are behind the paywall, ensure that the URL is in this format: `https://<username>:<password>@archive.cloudera.com/...`

For example, `https://<username>:<password>@archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/`

| Operating System | Version Number         | Repository Name | Format  | URL   |
|------------------|------------------------|-----------------|---|---|
| CentOS           | HDP-7.1.7.2000-305.xml | HDPDC           | Version Definition File (VDF)   | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/centos7/HDP-7.1.7.2000-305.xml">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/centos7/HDP-7.1.7.2000-305.xml</a>                                       |
|                  |                        |                 | Base URL  | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/centos7">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/centos7</a>   |
|                  |                        |                 | Repo File   | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/centos7/hdp.repo">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/centos7/hdp.repo</a>   |
|                  |                        |                 | Tarball   | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/centos7/HDP-7.1.7.2000-centos7-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/centos7/HDP-7.1.7.2000-centos7-rpm.tar.gz</a>                 |
|                  |                        | md5             | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/centos7/HDP-7.1.7.2000-centos7-rpm.tar.gz.md5">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/centos7/HDP-7.1.7.2000-centos7-rpm.tar.gz.md5</a> |   |
|                  |                        | asc             | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/centos7/HDP-7.1.7.2000-centos7-rpm.tar.gz.asc">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/centos7/HDP-7.1.7.2000-centos7-rpm.tar.gz.asc</a> |   |
|                  |                        | HDP-UTILS       | Base URL  | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/</a>   |
|                  |                        |                 | Tarball   | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz</a>         |
|                  |                        |                 | md5   | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.md5">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.md5</a> |
|                  |                        |                 | asc   | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.asc">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.asc</a> |

| Operating System | Version Number | Repository Name | Format  | URL   |
|------------------|----------------|-----------------|---------|---|
|                  |                | HDP-GPL         | URL     | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/centos7/hdp.gpl.repo">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/centos7/hdp.gpl.repo</a>   |
|                  |                |                 | Tarball | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/centos7/HDP-GPL-7.1.7.2000-centos7-gpl.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/centos7/HDP-GPL-7.1.7.2000-centos7-gpl.tar.gz</a>         |
|                  |                |                 | md5     | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/centos7/HDP-GPL-7.1.7.2000-centos7-gpl.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/centos7/HDP-GPL-7.1.7.2000-centos7-gpl.tar.gz.md5</a> |
|                  |                |                 | asc     | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/centos7/HDP-GPL-7.1.7.2000-centos7-gpl.tar.gz.asc">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/centos7/HDP-GPL-7.1.7.2000-centos7-gpl.tar.gz.asc</a> |

| Operating System | Version Number         | Repository Name | Format                        | URL   |
|------------------|------------------------|-----------------|-------------------------------|---|
| Ubuntu 18        | HDP-7.1.7.2000-305.xml | HDPDC           | Version Definition File (VDF) | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/ubuntu18/HDP-7.1.7.2000-305.xml">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/ubuntu18/HDP-7.1.7.2000-305.xml</a>                                 |
|                  |                        |                 | Base URL                      | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/ubuntu18">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/ubuntu18</a>   |
|                  |                        |                 | Repo File                     | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/ubuntu18/hdp.list">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/ubuntu18/hdp.list</a>   |
|                  |                        |                 | Tarball                       | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/ubuntu18/HDP-7.1.7.2000-ubuntu18-deb.tar.gz">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/ubuntu18/HDP-7.1.7.2000-ubuntu18-deb.tar.gz</a>         |
|                  |                        |                 | md5                           | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/ubuntu18/HDP-7.1.7.2000-ubuntu18-deb.tar.gz.md5">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/ubuntu18/HDP-7.1.7.2000-ubuntu18-deb.tar.gz.md5</a> |

| Operating System | Version Number | Repository Name | Format   | URL   |
|------------------|----------------|-----------------|----------|---|
|                  |                |                 |          | asc<br><a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/ubuntu18/HDP-7.1.7.2000-ubuntu18-deb.tar.gz.asc">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/ubuntu18/HDP-7.1.7.2000-ubuntu18-deb.tar.gz.asc</a>          |
|                  |                | HDP-UTILS       | Base URL | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/</a>   |
|                  |                |                 | Tarball  | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz</a>                 |
|                  |                |                 | md5      | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.md5">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.md5</a>         |
|                  |                |                 | asc      | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.asc">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.asc</a>         |
|                  |                | HDPDC-GPL       | URL      | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/ubuntu18/hdp.gpl.list">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/ubuntu18/hdp.gpl.list</a>   |
|                  |                |                 | Tarball  | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/ubuntu18/HDP-GPL-7.1.7.2000-ubuntu18-gpl.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/ubuntu18/HDP-GPL-7.1.7.2000-ubuntu18-gpl.tar.gz</a>         |
|                  |                |                 | md5      | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/ubuntu18/HDP-GPL-7.1.7.2000-ubuntu18-gpl.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/ubuntu18/HDP-GPL-7.1.7.2000-ubuntu18-gpl.tar.gz.md5</a> |

| Operating System | Version Number | Repository Name | Format | URL  |
|------------------|----------------|-----------------|--------|--|
|                  |                |                 |        | asc<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/ubuntu18/HDP-GPL-7.1.7.2000-ubuntu18-gpl.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/ubuntu18/HDP-GPL-7.1.7.2000-ubuntu18-gpl.tar.gz</a> |

| Operating System | Version Number         | Repository Name | Format                        | URL  |          |   |
|------------------|------------------------|-----------------|-------------------------------|--|----------|---|
| SLES 12          | HDP-7.1.7.2000-305.xml | HDPDC           | Version Definition File (VDF) | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/sles12/HDP-7.1.7.2000-305.xml">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/sles12/HDP-7.1.7.2000-305.xml</a>                                    |          |   |
|                  |                        |                 | Base URL                      | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/sles12">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/sles12</a>  |          |   |
|                  |                        |                 | Repo File                     | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/sles12/hdp.repo">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/sles12/hdp.repo</a>  |          |   |
|                  |                        |                 | Tarball                       | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/sles12/HDP-7.1.7.2000-sles12-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/sles12/HDP-7.1.7.2000-sles12-rpm.tar.gz</a>                |          |   |
|                  |                        |                 | md5                           | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/sles12/HDP-7.1.7.2000-sles12-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/sles12/HDP-7.1.7.2000-sles12-rpm.tar.gz</a>                |          |   |
|                  |                        |                 | asc                           | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/sles12/HDP-7.1.7.2000-sles12-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.2000/sles12/HDP-7.1.7.2000-sles12-rpm.tar.gz</a>                |          |   |
|                  |                        |                 |                               |  |          |   |
|                  |                        |                 |                               |  |          |   |
|                  |                        |                 |                               | HDP-UTILS  | Base URL | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12</a>   |
|                  |                        |                 |                               |  | Tarball  | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz</a> |
|                  |                        |                 |                               | md5<br><a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz</a> |          |   |

| Operating System | Version Number | Repository Name | Format  | URL  |
|------------------|----------------|-----------------|---------|--|
|                  |                |                 |         | asc<br><a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz.asc">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz.asc</a>         |
|                  |                | HDP-GPL         | URL     | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/sles12/hdp.gpl.repo">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/sles12/hdp.gpl.repo</a>  |
|                  |                |                 | Tarball | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/sles12/HDP-GPL-7.1.7.2000-sles12-gpl.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/sles12/HDP-GPL-7.1.7.2000-sles12-gpl.tar.gz</a>                |
|                  |                |                 |         | md5<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/sles12/HDP-GPL-7.1.7.2000-sles12-gpl.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/sles12/HDP-GPL-7.1.7.2000-sles12-gpl.tar.gz.md5</a> |
|                  |                |                 |         | asc<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/sles12/HDP-GPL-7.1.7.2000-sles12-gpl.tar.gz.asc">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.2000/sles12/HDP-GPL-7.1.7.2000-sles12-gpl.tar.gz.asc</a> |

| Operating System | Version Number      | Repository Name | Format                        | URL   |
|------------------|---------------------|-----------------|-------------------------------|---|
| CentOS           | HDP-7.1.7.78-12.xml | HDPDC           | Version Definition File (VDF) | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/centos7/HDP-7.1.7.78-12.xml">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/centos7/HDP-7.1.7.78-12.xml</a>                         |
|                  |                     |                 | Base URL                      | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/centos7">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/centos7</a>   |
|                  |                     |                 | Repo File                     | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/centos7/hdp.repo">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/centos7/hdp.repo</a>   |
|                  |                     |                 | Tarball                       | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/centos7/HDP-7.1.7.78-centos7-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/centos7/HDP-7.1.7.78-centos7-rpm.tar.gz</a> |

| Operating System | Version Number | Repository Name | Format   | URL  |
|------------------|----------------|-----------------|----------|--|
|                  |                |                 |          | md5<br><a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/centos7/HDP-7.1.7.78-centos7-rpm.tar.gz.md5">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/centos7/HDP-7.1.7.78-centos7-rpm.tar.gz.md5</a>                 |
|                  |                |                 |          | asc<br><a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/centos7/HDP-7.1.7.78-centos7-rpm.tar.gz.asc">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/centos7/HDP-7.1.7.78-centos7-rpm.tar.gz.asc</a>                 |
|                  |                | HDP-UTILS       | Base URL | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/</a>  |
|                  |                |                 | Tarball  | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz</a>                |
|                  |                |                 |          | md5<br><a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.md5">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.md5</a> |
|                  |                |                 |          | asc<br><a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.asc">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.asc</a> |
|                  |                | HDP-GPL         | URL      | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/centos7/hdp.gpl.repo">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/centos7/hdp.gpl.repo</a>  |
|                  |                |                 |          | Tarball  |
|                  |                |                 |          | md5<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/centos7/HDP-GPL-7.1.7.78-centos7-gpl.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/centos7/HDP-GPL-7.1.7.78-centos7-gpl.tar.gz.md5</a> |

| Operating System | Version Number | Repository Name | Format | URL  |
|------------------|----------------|-----------------|--------|--|
|                  |                |                 |        | asc<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/centos7/HDP-GPL-7.1.7.78-centos7-gpl.tar.gz.asc">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/centos7/HDP-GPL-7.1.7.78-centos7-gpl.tar.gz.asc</a> |

| Operating System | Version Number      | Repository Name | Format                        | URL   |   |   |
|------------------|---------------------|-----------------|-------------------------------|---|---|---|
| Ubuntu 18        | HDP-7.1.7.78-12.xml | HDPDC           | Version Definition File (VDF) | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/ubuntu18/HDP-7.1.7.78-12.xml">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/ubuntu18/HDP-7.1.7.78-12.xml</a>                                   |   |   |
|                  |                     |                 | Base URL                      | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/ubuntu18">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/ubuntu18</a>   |   |   |
|                  |                     |                 | Repo File                     | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/ubuntu18/hdp.list">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/ubuntu18/hdp.list</a>   |   |   |
|                  |                     |                 | Tarball                       | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/ubuntu18/HDP-7.1.7.78-ubuntu18-deb.tar.gz">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/ubuntu18/HDP-7.1.7.78-ubuntu18-deb.tar.gz</a>         |   |   |
|                  |                     |                 | md5                           | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/ubuntu18/HDP-7.1.7.78-ubuntu18-deb.tar.gz.md5">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/ubuntu18/HDP-7.1.7.78-ubuntu18-deb.tar.gz.md5</a> |   |   |
|                  |                     |                 | asc                           | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/ubuntu18/HDP-7.1.7.78-ubuntu18-deb.tar.gz.asc">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/ubuntu18/HDP-7.1.7.78-ubuntu18-deb.tar.gz.asc</a> |   |   |
|                  |                     |                 |                               |   |   |   |
|                  |                     |                 |                               |   |   |   |
|                  |                     |                 |                               | HDP-UTILS   | Base URL  | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/</a>   |
|                  |                     |                 |                               |   | Tarball   | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz</a> |
|                  |                     |                 |                               | md5   | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.md5">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.md5</a> |   |



| Operating System | Version Number | Repository Name | Format  | URL  |
|------------------|----------------|-----------------|---------|--|
|                  |                |                 |         | asc<br><a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz</a>         |
|                  |                | HDPDC-GPL       | URL     | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/ubuntu18/hdp.gpl.list">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/ubuntu18/hdp.gpl.list</a>  |
|                  |                |                 | Tarball | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/ubuntu18/HDP-GPL-7.1.7.78-ubuntu18-gpl.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/ubuntu18/HDP-GPL-7.1.7.78-ubuntu18-gpl.tar.gz</a>                |
|                  |                |                 |         | md5<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/ubuntu18/HDP-GPL-7.1.7.78-ubuntu18-gpl.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/ubuntu18/HDP-GPL-7.1.7.78-ubuntu18-gpl.tar.gz.md5</a> |
|                  |                |                 |         | asc<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/ubuntu18/HDP-GPL-7.1.7.78-ubuntu18-gpl.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/ubuntu18/HDP-GPL-7.1.7.78-ubuntu18-gpl.tar.gz</a>         |

| Operating System | Version Number      | Repository Name | Format                        | URL  |
|------------------|---------------------|-----------------|-------------------------------|--|
| SLES 12          | HDP-7.1.7.78-12.xml | HDPDC           | Version Definition File (VDF) | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/sles12/HDP-7.1.7.78-12.xml">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/sles12/HDP-7.1.7.78-12.xml</a>                              |
|                  |                     |                 | Base URL                      | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/sles12">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/sles12</a>  |
|                  |                     |                 | Repo File                     | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/sles12/hdp.repo">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/sles12/hdp.repo</a>  |
|                  |                     |                 | Tarball                       | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/sles12/HDP-7.1.7.78-sles12-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/sles12/HDP-7.1.7.78-sles12-rpm.tar.gz</a>        |
|                  |                     |                 |                               | md5<br><a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/sles12/HDP-7.1.7.78-sles12-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/sles12/HDP-7.1.7.78-sles12-rpm.tar.gz</a> |

| Operating System | Version Number | Repository Name | Format   | URL   |
|------------------|----------------|-----------------|----------|---|
|                  |                |                 |          | asc<br><a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/sles12/HDP-7.1.7.78-sles12-rpm.tar.gz.asc">https://archive.cloudera.com/p/HDPDC/7.x/7.1.7.78/sles12/HDP-7.1.7.78-sles12-rpm.tar.gz.asc</a>          |
|                  |                | HDP-UTILS       | Base URL | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12</a>   |
|                  |                |                 | Tarball  | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz</a>         |
|                  |                |                 | md5      | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz.md5">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz.md5</a> |
|                  |                |                 | asc      | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz.asc">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz.asc</a> |
|                  |                | HDP-GPL         | URL      | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/sles12/hdp.gpl.repo">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/sles12/hdp.gpl.repo</a>   |
|                  |                |                 | Tarball  | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/sles12/HDP-GPL-7.1.7.78-sles12-gpl.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/sles12/HDP-GPL-7.1.7.78-sles12-gpl.tar.gz</a>         |
|                  |                |                 | md5      | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/sles12/HDP-GPL-7.1.7.78-sles12-gpl.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/sles12/HDP-GPL-7.1.7.78-sles12-gpl.tar.gz.md5</a> |
|                  |                |                 | asc      | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/sles12/HDP-GPL-7.1.7.78-sles12-gpl.tar.gz.asc">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.7.78/sles12/HDP-GPL-7.1.7.78-sles12-gpl.tar.gz.asc</a> |

| Operating System | Version Number      | Repository Name | Format                        | URL   |   |
|------------------|---------------------|-----------------|-------------------------------|---|---|
| CentOS           | HDP-7.1.6.0-203.xml | HDPDC           | Version Definition File (VDF) | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/centos7/HDP-7.1.6.0-297.xml">https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/centos7/HDP-7.1.6.0-297.xml</a>                                   |   |
|                  |                     |                 | Base URL                      | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/centos7">https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/centos7</a>   |   |
|                  |                     |                 | Repo File                     | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/centos7/hdpdc.repo">https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/centos7/hdpdc.repo</a>   |   |
|                  |                     |                 | Tarball                       | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/centos7/HDPDC-7.1.6.0-centos7-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/centos7/HDPDC-7.1.6.0-centos7-rpm.tar.gz</a>         |   |
|                  |                     |                 | md5                           | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/centos7/HDPDC-7.1.6.0-centos7-rpm.tar.gz.md5">https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/centos7/HDPDC-7.1.6.0-centos7-rpm.tar.gz.md5</a> |   |
|                  |                     |                 | asc                           | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/centos7/HDPDC-7.1.6.0-centos7-rpm.tar.gz.asc">https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/centos7/HDPDC-7.1.6.0-centos7-rpm.tar.gz.asc</a> |   |
|                  |                     |                 | HDP-UTILS                     | Base URL  | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/</a>   |
|                  |                     |                 |                               | Tarball   | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz</a>         |
|                  |                     |                 |                               | md5   | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.md5">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.md5</a> |
|                  |                     |                 |                               | asc   | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.asc">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.asc</a> |

| Operating System | Version Number | Repository Name | Format  | URL  |
|------------------|----------------|-----------------|---------|--|
|                  |                | HDP-GPL         | URL     | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/centos7/hdpdc.gpl.repo">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/centos7/hdpdc.gpl.repo</a>  |
|                  |                |                 | Tarball | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/centos7/HDPDC-GPL-7.1.6.0-centos7-tars-tarball.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/centos7/HDPDC-GPL-7.1.6.0-centos7-tars-tarball.tar.gz</a>                |
|                  |                |                 |         | md5<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/centos7/HDPDC-GPL-7.1.6.0-centos7-tars-tarball.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/centos7/HDPDC-GPL-7.1.6.0-centos7-tars-tarball.tar.gz.md5</a> |
|                  |                |                 |         | asc<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/centos7/HDPDC-GPL-7.1.6.0-centos7-tars-tarball.tar.gz.asc">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/centos7/HDPDC-GPL-7.1.6.0-centos7-tars-tarball.tar.gz.asc</a> |

| Operating System   | Version Number      | Repository Name | Format                        | URL   |
|--|---------------------|-----------------|-------------------------------|---|
| Ubuntu 18  | HDP-7.1.6.0-203.xml | HDPDC           | Version Definition File (VDF) | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/ubuntu18/HDP-7.1.6.0-297.xml">https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/ubuntu18/HDP-7.1.6.0-297.xml</a>                             |
|  |                     |                 | Base URL                      | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/ubuntu18/7.1.6.0 HDP main">https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/ubuntu18/7.1.6.0 HDP main</a>                                   |
|  |                     |                 | Repo File                     | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/ubuntu18/hdpdc.list">https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/ubuntu18/hdpdc.list</a>   |
|  |                     |                 | Tarball                       | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/ubuntu18/HDPDC-7.1.6.0-ubuntu18-deb.tar.gz">https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/ubuntu18/HDPDC-7.1.6.0-ubuntu18-deb.tar.gz</a> |
| md5<br><a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/ubuntu18/HDPDC-7.1.6.0-ubuntu18-deb.tar.gz.md5">https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/ubuntu18/HDPDC-7.1.6.0-ubuntu18-deb.tar.gz.md5</a> |                     |                 |                               |   |
| asc<br><a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/ubuntu18/HDPDC-7.1.6.0-ubuntu18-deb.tar.gz.asc">https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/ubuntu18/HDPDC-7.1.6.0-ubuntu18-deb.tar.gz.asc</a> |                     |                 |                               |   |

| Operating System | Version Number | Repository Name | Format  | URL   |   |
|------------------|----------------|-----------------|---|---|---|
|                  |                | HDP-UTILS       | Base URL  | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/</a>   |   |
|                  |                |                 | Tarball   | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz</a> |   |
|                  |                |                 |   | md5   | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.md5">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.md5</a> |
|                  |                |                 |   | asc   | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.asc">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.asc</a> |
|                  |                |                 | HDPDC-GPL   | URL   | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/ubuntu18/hdpdc.gpl.list">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/ubuntu18/hdpdc.gpl.list</a>   |
|                  |                |                 |   | Tarball   | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/ubuntu18/HDPDC-GPL-7.1.6.0-ubuntu18-deb.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/ubuntu18/HDPDC-GPL-7.1.6.0-ubuntu18-deb.tar.gz</a>         |
|                  |                | md5             |   |   | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/ubuntu18/HDPDC-GPL-7.1.6.0-ubuntu18-deb.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/ubuntu18/HDPDC-GPL-7.1.6.0-ubuntu18-deb.tar.gz.md5</a> |
|                  |                | asc             | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/ubuntu18/HDPDC-GPL-7.1.6.0-ubuntu18-deb.tar.gz.asc">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/ubuntu18/HDPDC-GPL-7.1.6.0-ubuntu18-deb.tar.gz.asc</a> |   |   |

| Operating System | Version Number      | Repository Name | Format                        | URL   |
|------------------|---------------------|-----------------|-------------------------------|---|
| SLES 12          | HDP-7.1.6.0-203.xml | HDPDC           | Version Definition File (VDF) | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/sles12/HDP-7.1.6.0-297.xml">https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/sles12/HDP-7.1.6.0-297.xml</a> |

| Operating System | Version Number | Repository Name | Format    | URL   |
|------------------|----------------|-----------------|-----------|---|
|                  |                |                 | Base URL  | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/sles12">https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/sles12</a>   |
|                  |                |                 | Repo File | <a href="https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/sles12/hdpdc.repo">https://archive.cloudera.com/p/HDPDC/7.x/7.1.6.0/sles12/hdpdc.repo</a>   |
|                  |                |                 | Tarball   | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/sles12/HDPDC-GPL-7.1.6.0-sles12-tars-tarball.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/sles12/HDPDC-GPL-7.1.6.0-sles12-tars-tarball.tar.gz</a>         |
|                  |                |                 | md5       | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/sles12/HDPDC-GPL-7.1.6.0-sles12-tars-tarball.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/sles12/HDPDC-GPL-7.1.6.0-sles12-tars-tarball.tar.gz.md5</a> |
|                  |                |                 | asc       | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/sles12/HDPDC-GPL-7.1.6.0-sles12-tars-tarball.tar.gz.asc">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/sles12/HDPDC-GPL-7.1.6.0-sles12-tars-tarball.tar.gz.asc</a> |
|                  |                | HDP-UTILS       | Base URL  | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12</a>   |
|                  |                |                 | Tarball   | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz</a>                           |
|                  |                |                 | md5       | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz.md5">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz.md5</a>                   |
|                  |                |                 | asc       | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz.asc">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz.asc</a>                   |
|                  |                | HDP-GPL         | URL       | <a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/sles12/hdpdc.gpl.repo">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/sles12/hdpdc.gpl.repo</a>   |

| Operating System | Version Number | Repository Name | Format  | URL   |
|------------------|----------------|-----------------|---------|---|
|                  |                |                 | Tarball | <a href="https://archive.cloudera.com/p/HDPDC-GPL/sles12/7.x/updates/7.1.6.0/HDPDC-GPL-7.1.6.0-sles12-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/sles12/7.x/updates/7.1.6.0/HDPDC-GPL-7.1.6.0-sles12-rpm.tar.gz</a> |
|                  |                |                 |         | md5<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/sles12/HDPDC-GPL-7.1.6.0-sles12-rpm.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/sles12/HDPDC-GPL-7.1.6.0-sles12-rpm.tar.gz.md5</a>  |
|                  |                |                 |         | asc<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/sles12/HDPDC-GPL-7.1.6.0-sles12-rpm.tar.gz.asc">https://archive.cloudera.com/p/HDPDC-GPL/7.x/7.1.6.0/sles12/HDPDC-GPL-7.1.6.0-sles12-rpm.tar.gz.asc</a>  |

| Operating System | Version Number      | Repository Name | Format                        | URL  |
|------------------|---------------------|-----------------|-------------------------------|--|
| CentOS           | HDP-7.1.4.0-203.xml | HDPDC           | Version Definition File (VDF) | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.4.0/HDP-7.1.4.0-203.xml">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.4.0/HDP-7.1.4.0-203.xml</a>  |
|                  |                     |                 | Base URL                      | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.4.0">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.4.0</a>  |
|                  |                     |                 | Repo File                     | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.4.0/hdpdc.repo">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.4.0/hdpdc.repo</a>  |
|                  |                     |                 | Tarball                       | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.4.0/HDPDC-7.1.4.0-centos7-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.4.0/HDPDC-7.1.4.0-centos7-rpm.tar.gz</a>                |
|                  |                     |                 |                               | md5<br><a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.4.0/HDPDC-7.1.4.0-centos7-rpm.tar.gz.md5">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.4.0/HDPDC-7.1.4.0-centos7-rpm.tar.gz.md5</a> |
|                  |                     |                 |                               | asc<br><a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.4.0/HDPDC-7.1.4.0-centos7-rpm.tar.gz.asc">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.4.0/HDPDC-7.1.4.0-centos7-rpm.tar.gz.asc</a> |

| Operating System | Version Number      | Repository Name | Format                        | URL  |
|------------------|---------------------|-----------------|-------------------------------|--|
|                  |                     | HDP-UTILS       | Base URL                      | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/</a>  |
|                  |                     |                 | Tarball                       | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz</a>                                |
|                  |                     |                 |                               | md5<br><a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.md5">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.md5</a>                 |
|                  |                     |                 |                               | asc<br><a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.asc">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.asc</a>                 |
|                  |                     | HDP-GPL         | URL                           | <a href="https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.4.0/hdpdc.gpl.repo">https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.4.0/hdpdc.gpl.repo</a>  |
|                  |                     |                 | Tarball                       | <a href="https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.4.0/HDPDC-GPL-7.1.4.0-centos7-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.4.0/HDPDC-GPL-7.1.4.0-centos7-rpm.tar.gz</a>                |
|                  |                     |                 |                               | md5<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.4.0/HDPDC-GPL-7.1.4.0-centos7-rpm.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.4.0/HDPDC-GPL-7.1.4.0-centos7-rpm.tar.gz.md5</a> |
|                  |                     |                 |                               | asc<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.4.0/HDPDC-GPL-7.1.4.0-centos7-rpm.tar.gz.asc">https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.4.0/HDPDC-GPL-7.1.4.0-centos7-rpm.tar.gz.asc</a> |
| Operating System | Version Number      | Repository Name | Format                        | URL  |
| Ubuntu 18        | HDP-7.1.4.0-203.xml | HDPDC           | Version Definition File (VDF) | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.4.0/HDP-7.1.4.0-203.xml">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.4.0/HDP-7.1.4.0-203.xml</a>  |



| Operating System | Version Number | Repository Name | Format    | URL   |
|------------------|----------------|-----------------|-----------|---|
|                  |                |                 | Base URL  | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.4.0/HDP%20main">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.4.0/HDP main</a>   |
|                  |                |                 | Repo File | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.4.0/hdpdc.list">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.4.0/hdpdc.list</a>   |
|                  |                |                 | Tarball   | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.4.0/HDPDC-7.1.4.0-ubuntu18-deb.tar.gz">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.4.0/HDPDC-7.1.4.0-ubuntu18-deb.tar.gz</a>         |
|                  |                |                 | md5       | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.4.0/HDPDC-7.1.4.0-ubuntu18-deb.tar.gz.md5">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.4.0/HDPDC-7.1.4.0-ubuntu18-deb.tar.gz.md5</a> |
|                  |                |                 | asc       | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.4.0/HDPDC-7.1.4.0-ubuntu18-deb.tar.gz.asc">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.4.0/HDPDC-7.1.4.0-ubuntu18-deb.tar.gz.asc</a> |
|                  |                | HDP-UTILS       | Base URL  | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/</a>   |
|                  |                |                 | Tarball   | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz</a>         |
|                  |                |                 | md5       | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.md5">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.md5</a> |
|                  |                |                 | asc       | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.asc">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.asc</a> |

| Operating System | Version Number | Repository Name | Format  | URL   |
|------------------|----------------|-----------------|---------|---|
|                  |                | HDPDC-GPL       | URL     | <a href="https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.4.0/hdpdc.gpl.list">https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.4.0/hdpdc.gpl.list</a>   |
|                  |                |                 | Tarball | <a href="https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.4.0/HDPDC-GPL-7.1.4.0-ubuntu18-deb.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.4.0/HDPDC-GPL-7.1.4.0-ubuntu18-deb.tar.gz</a>         |
|                  |                |                 | md5     | <a href="https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.4.0/HDPDC-GPL-7.1.4.0-ubuntu18-deb.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.4.0/HDPDC-GPL-7.1.4.0-ubuntu18-deb.tar.gz.md5</a> |
|                  |                |                 | asc     | <a href="https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.4.0/HDPDC-GPL-7.1.4.0-ubuntu18-deb.tar.gz.asc">https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.4.0/HDPDC-GPL-7.1.4.0-ubuntu18-deb.tar.gz.asc</a> |

| Operating System | Version Number      | Repository Name | Format                        | URL   |
|------------------|---------------------|-----------------|-------------------------------|---|
| SLES 12          | HDP-7.1.4.0-203.xml | HDPDC           | Version Definition File (VDF) | <a href="https://archive.cloudera.com/p/HDPDC/sles12/7.x/updates/7.1.4.0/HDP-7.1.4.0-203.xml">https://archive.cloudera.com/p/HDPDC/sles12/7.x/updates/7.1.4.0/HDP-7.1.4.0-203.xml</a>                                 |
|                  |                     |                 | Base URL                      | <a href="https://archive.cloudera.com/p/HDPDC/sles12/7.x/updates/7.1.4.0">https://archive.cloudera.com/p/HDPDC/sles12/7.x/updates/7.1.4.0</a>   |
|                  |                     |                 | Repo File                     | <a href="https://archive.cloudera.com/p/HDPDC/sles12/7.x/updates/7.1.4.0/hdpdc.repo">https://archive.cloudera.com/p/HDPDC/sles12/7.x/updates/7.1.4.0/hdpdc.repo</a>   |
|                  |                     |                 | Tarball                       | <a href="https://archive.cloudera.com/p/HDPDC/sles12/7.x/updates/7.1.4.0/HDPDC-7.1.4.0-sles12-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC/sles12/7.x/updates/7.1.4.0/HDPDC-7.1.4.0-sles12-rpm.tar.gz</a>         |
|                  |                     |                 | md5                           | <a href="https://archive.cloudera.com/p/HDPDC/sles12/7.x/updates/7.1.4.0/HDPDC-7.1.4.0-sles12-rpm.tar.gz.md5">https://archive.cloudera.com/p/HDPDC/sles12/7.x/updates/7.1.4.0/HDPDC-7.1.4.0-sles12-rpm.tar.gz.md5</a> |

| Operating System | Version Number | Repository Name | Format   | URL   |
|------------------|----------------|-----------------|----------|---|
|                  |                |                 |          | asc<br><a href="https://archive.cloudera.com/p/HDPDC/sles12/7.x/updates/7.1.4.0/HDPDC-7.1.4.0-sles12-rpm.tar.gz.asc">https://archive.cloudera.com/p/HDPDC/sles12/7.x/updates/7.1.4.0/HDPDC-7.1.4.0-sles12-rpm.tar.gz.asc</a>          |
|                  |                | HDP-UTILS       | Base URL | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/</a>   |
|                  |                |                 | Tarball  | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz</a>                         |
|                  |                |                 | md5      | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz.md5">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz.md5</a>                 |
|                  |                |                 | asc      | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz.asc">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/sles12/HDP-UTILS-1.1.0.22-sles12.tar.gz.asc</a>                 |
|                  |                | HDP-GPL         | URL      | <a href="https://archive.cloudera.com/p/HDPDC-GPL/sles12/7.x/updates/7.1.4.0/hdpdc.gpl.repo">https://archive.cloudera.com/p/HDPDC-GPL/sles12/7.x/updates/7.1.4.0/hdpdc.gpl.repo</a>   |
|                  |                |                 | Tarball  | <a href="https://archive.cloudera.com/p/HDPDC-GPL/sles12/7.x/updates/7.1.4.0/HDPDC-GPL-7.1.4.0-sles12-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/sles12/7.x/updates/7.1.4.0/HDPDC-GPL-7.1.4.0-sles12-rpm.tar.gz</a>         |
|                  |                |                 | md5      | <a href="https://archive.cloudera.com/p/HDPDC-GPL/sles12/7.x/updates/7.1.4.0/HDPDC-GPL-7.1.4.0-sles12-rpm.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/sles12/7.x/updates/7.1.4.0/HDPDC-GPL-7.1.4.0-sles12-rpm.tar.gz.md5</a> |
|                  |                |                 | asc      | <a href="https://archive.cloudera.com/p/HDPDC-GPL/sles12/7.x/updates/7.1.4.0/HDPDC-GPL-7.1.4.0-sles12-rpm.tar.gz.asc">https://archive.cloudera.com/p/HDPDC-GPL/sles12/7.x/updates/7.1.4.0/HDPDC-GPL-7.1.4.0-sles12-rpm.tar.gz.asc</a> |

| Operating System | Version Number      | Repository Name | Format  | URL   |
|------------------|---------------------|-----------------|---|---|
| CentOS           | HDP-7.1.3.0-100.xml | HDPDC           | Version Definition File (VDF)   | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.3.0/HDP-7.1.3.0-100.xml">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.3.0/HDP-7.1.3.0-100.xml</a>                           |
|                  |                     |                 | Base URL  | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.3.0">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.3.0</a>   |
|                  |                     |                 | Repo File   | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.3.0/hdpdc.repo">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.3.0/hdpdc.repo</a>   |
|                  |                     |                 | Tarball   | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.3.0/HDPDC-7.1.3.0-centos7-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.3.0/HDPDC-7.1.3.0-centos7-rpm.tar.gz</a> |
|                  |                     | md5             | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.3.0/HDPDC-7.1.3.0-centos7-rpm.tar.gz.md5">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.3.0/HDPDC-7.1.3.0-centos7-rpm.tar.gz.md5</a> |   |
|                  |                     | asc             | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.3.0/HDPDC-7.1.3.0-centos7-rpm.tar.gz.asc">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.3.0/HDPDC-7.1.3.0-centos7-rpm.tar.gz.asc</a> |   |
|                  |                     | HDP-UTILS       | Base URL  | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/</a>   |
|                  |                     | Tarball         | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz</a>         |   |
|                  |                     | md5             | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.md5">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.md5</a> |   |
|                  |                     | asc             | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.asc">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.asc</a> |   |

| Operating System | Version Number | Repository Name | Format  | URL   |
|------------------|----------------|-----------------|---------|---|
|                  |                | HDP-GPL         | URL     | <a href="https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.3.0/hdpdc.gpl.repo">https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.3.0/hdpdc.gpl.repo</a>   |
|                  |                |                 | Tarball | <a href="https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.3.0/HDPDC-GPL-7.1.3.0-centos7-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.3.0/HDPDC-GPL-7.1.3.0-centos7-rpm.tar.gz</a>         |
|                  |                |                 | md5     | <a href="https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.3.0/HDPDC-GPL-7.1.3.0-centos7-rpm.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.3.0/HDPDC-GPL-7.1.3.0-centos7-rpm.tar.gz.md5</a> |
|                  |                |                 | asc     | <a href="https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.3.0/HDPDC-GPL-7.1.3.0-centos7-rpm.tar.gz.asc">https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.3.0/HDPDC-GPL-7.1.3.0-centos7-rpm.tar.gz.asc</a> |

| Operating System | Version Number      | Repository Name | Format                        | URL   |
|------------------|---------------------|-----------------|-------------------------------|---|
| Ubuntu 18        | HDP-7.1.3.0-100.xml | HDPDC           | Version Definition File (VDF) | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.3.0/HDP-7.1.3.0-100.xml">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.3.0/HDP-7.1.3.0-100.xml</a>                                     |
|                  |                     |                 | Base URL                      | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.3.0/HDP main">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.3.0/HDP main</a>   |
|                  |                     |                 | Repo File                     | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.3.0/hdpdc.list">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.3.0/hdpdc.list</a>   |
|                  |                     |                 | Tarball                       | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.3.0/HDPDC-7.1.3.0-ubuntu18-deb.tar.gz">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.3.0/HDPDC-7.1.3.0-ubuntu18-deb.tar.gz</a>         |
|                  |                     |                 | md5                           | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.3.0/HDPDC-7.1.3.0-ubuntu18-deb.tar.gz.md5">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.3.0/HDPDC-7.1.3.0-ubuntu18-deb.tar.gz.md5</a> |

| Operating System | Version Number | Repository Name | Format   | URL   |
|------------------|----------------|-----------------|----------|---|
|                  |                |                 |          | asc<br><a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.3.0/HDPDC-7.1.3.0-ubuntu18-deb.tar.gz.asc">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.3.0/HDPDC-7.1.3.0-ubuntu18-deb.tar.gz.asc</a>          |
|                  |                | HDP-UTILS       | Base URL | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/</a>   |
|                  |                |                 | Tarball  | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz</a>                         |
|                  |                |                 | md5      | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.md5">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.md5</a>                 |
|                  |                |                 | asc      | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.asc">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.asc</a>                 |
|                  |                | HDPDC-GPL       | URL      | <a href="https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.3.0/hdpdc.gpl.list">https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.3.0/hdpdc.gpl.list</a>   |
|                  |                |                 | Tarball  | <a href="https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.3.0/HDPDC-GPL-7.1.3.0-ubuntu18-deb.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.3.0/HDPDC-GPL-7.1.3.0-ubuntu18-deb.tar.gz</a>         |
|                  |                |                 | md5      | <a href="https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.3.0/HDPDC-GPL-7.1.3.0-ubuntu18-deb.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.3.0/HDPDC-GPL-7.1.3.0-ubuntu18-deb.tar.gz.md5</a> |

| Operating System | Version Number | Repository Name | Format | URL  |
|------------------|----------------|-----------------|--------|--|
|                  |                |                 |        | asc<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.3.0/HDPDC-GPL-7.1.3.0-ubuntu18-deb.tar.gz.asc">https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.3.0/HDPDC-GPL-7.1.3.0-ubuntu18-deb.tar.gz.asc</a> |

| Operating System | Version Number     | Repository Name | Format                        | URL   |   |
|------------------|--------------------|-----------------|-------------------------------|---|---|
| CentOS           | HDP-7.1.2.0-96.xml | HDPDC           | Version Definition File (VDF) | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.2.0/HDP-7.1.2.0-96.xml">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.2.0/HDP-7.1.2.0-96.xml</a>                                     |   |
|                  |                    |                 | Base URL                      | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.2.0">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.2.0</a>   |   |
|                  |                    |                 | Repo File                     | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.2.0/hpdc.repo">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.2.0/hpdc.repo</a>   |   |
|                  |                    |                 | Tarball                       | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.2.0/HDPDC-7.1.2.0-centos7-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.2.0/HDPDC-7.1.2.0-centos7-rpm.tar.gz</a>         |   |
|                  |                    |                 | md5                           | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.2.0/HDPDC-7.1.2.0-centos7-rpm.tar.gz.md5">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.2.0/HDPDC-7.1.2.0-centos7-rpm.tar.gz.md5</a> |   |
|                  |                    |                 | asc                           | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.2.0/HDPDC-7.1.2.0-centos7-rpm.tar.gz.asc">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.2.0/HDPDC-7.1.2.0-centos7-rpm.tar.gz.asc</a> |   |
|                  |                    |                 | HDP-UTILS                     | Base URL  | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/</a> |
|                  |                    | Tarball         |                               | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz</a>         |   |

| Operating System | Version Number | Repository Name | Format  | URL  |
|------------------|----------------|-----------------|---------|--|
|                  |                |                 |         | md5<br><a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.md5">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.md5</a>                 |
|                  |                |                 |         | asc<br><a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.asc">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.asc</a>                 |
|                  |                | HDP-GPL         | URL     | <a href="https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.2.0/hdpdc.gpl.repo">https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.2.0/hdpdc.gpl.repo</a>  |
|                  |                |                 | Tarball | <a href="https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.2.0/HDPDC-GPL-7.1.2.0-centos7-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.2.0/HDPDC-GPL-7.1.2.0-centos7-rpm.tar.gz</a>                |
|                  |                |                 |         | md5<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.2.0/HDPDC-GPL-7.1.2.0-centos7-rpm.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.2.0/HDPDC-GPL-7.1.2.0-centos7-rpm.tar.gz.md5</a> |
|                  |                |                 |         | asc<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.2.0/HDPDC-GPL-7.1.2.0-centos7-rpm.tar.gz.asc">https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.2.0/HDPDC-GPL-7.1.2.0-centos7-rpm.tar.gz.asc</a> |

| Operating System | Version Number     | Repository Name | Format                        | URL   |
|------------------|--------------------|-----------------|-------------------------------|---|
| Ubuntu 18        | HDP-7.1.2.0-96.xml | HDPDC           | Version Definition File (VDF) | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.2.0/HDP-7.1.2.0-96.xml">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.2.0/HDP-7.1.2.0-96.xml</a> |
|                  |                    |                 | Base URL                      | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.2.0 HDP main">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.2.0 HDP main</a>                     |
|                  |                    |                 | Repo File                     | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.2.0/hdpdc.list">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.2.0/hdpdc.list</a>                 |



| Operating System | Version Number | Repository Name | Format   | URL   |
|------------------|----------------|-----------------|----------|---|
|                  |                |                 | Tarball  | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.2.0/HDPDC-7.1.2.0-ubuntu18-deb.tar.gz">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.2.0/HDPDC-7.1.2.0-ubuntu18-deb.tar.gz</a>                 |
|                  |                |                 |          | md5<br><a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.2.0/HDPDC-7.1.2.0-ubuntu18-deb.tar.gz.md5">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.2.0/HDPDC-7.1.2.0-ubuntu18-deb.tar.gz.md5</a>  |
|                  |                |                 |          | asc<br><a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.2.0/HDPDC-7.1.2.0-ubuntu18-deb.tar.gz.asc">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.2.0/HDPDC-7.1.2.0-ubuntu18-deb.tar.gz.asc</a>  |
|                  |                | HDP-UTILS       | Base URL | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/</a>   |
|                  |                |                 | Tarball  | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz</a>                 |
|                  |                |                 |          | md5<br><a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.md5">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.md5</a>  |
|                  |                |                 |          | asc<br><a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.asc">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.asc</a>  |
|                  |                | HDPDC-GPL       | URL      | <a href="https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.2.0/hdpdc.gpl.list">https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.2.0/hdpdc.gpl.list</a>   |
|                  |                |                 | Tarball  | <a href="https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.2.0/HDPDC-GPL-7.1.2.0-ubuntu18-deb.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.2.0/HDPDC-GPL-7.1.2.0-ubuntu18-deb.tar.gz</a> |

| Operating System | Version Number | Repository Name | Format | URL  |
|------------------|----------------|-----------------|--------|--|
|                  |                |                 |        | md5<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.2.0/HDPDC-GPL-7.1.2.0-ubuntu18-deb.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.2.0/HDPDC-GPL-7.1.2.0-ubuntu18-deb.tar.gz.md5</a> |
|                  |                |                 |        | asc<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.2.0/HDPDC-GPL-7.1.2.0-ubuntu18-deb.tar.gz.asc">https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.2.0/HDPDC-GPL-7.1.2.0-ubuntu18-deb.tar.gz.asc</a> |

| Operating System | Version Number  | Repository Name | Format   | URL   |
|------------------|-----------------|-----------------|--|---|
| CentOS           | HDP-7.1.1.0-565 | HDPDC           | Version Definition File (VDF)  | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.1.0/HDP-7.1.1.0-565.xml">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.1.0/HDP-7.1.1.0-565.xml</a>                           |
|                  |                 |                 | Base URL   | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.1.0">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.1.0</a>   |
|                  |                 |                 | Repo File  | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.1.0/hdpdc.repo">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.1.0/hdpdc.repo</a>   |
|                  |                 |                 | Tarball  | <a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.1.0/HDPDC-7.1.1.0-centos7-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.1.0/HDPDC-7.1.1.0-centos7-rpm.tar.gz</a> |
|                  |                 |                 | md5<br><a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.1.0/HDPDC-7.1.1.0-centos7-rpm.tar.gz.md5">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.1.0/HDPDC-7.1.1.0-centos7-rpm.tar.gz.md5</a> |   |
|                  |                 |                 | asc<br><a href="https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.1.0/HDPDC-7.1.1.0-centos7-rpm.tar.gz.asc">https://archive.cloudera.com/p/HDPDC/centos7/7.x/updates/7.1.1.0/HDPDC-7.1.1.0-centos7-rpm.tar.gz.asc</a> |   |
|                  |                 | HDP-UTILS       | Base URL   | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/</a>   |

| Operating System | Version Number | Repository Name | Format  | URL  |
|------------------|----------------|-----------------|---------|--|
|                  |                |                 | Tarball | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz</a>                                |
|                  |                |                 |         | md5<br><a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.md5">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.md5</a>                 |
|                  |                |                 |         | asc<br><a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.asc">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/centos7/HDP-UTILS-1.1.0.22-centos7.tar.gz.asc</a>                 |
|                  |                | HDP-GPL         | URL     | <a href="https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.1.0/hdpdc.gpl.repo">https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.1.0/hdpdc.gpl.repo</a>  |
|                  |                |                 | Tarball | <a href="https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.1.0/HDPDC-GPL-7.1.1.0-centos7-rpm.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.1.0/HDPDC-GPL-7.1.1.0-centos7-rpm.tar.gz</a>                |
|                  |                |                 |         | md5<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.1.0/HDPDC-GPL-7.1.1.0-centos7-rpm.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.1.0/HDPDC-GPL-7.1.1.0-centos7-rpm.tar.gz.md5</a> |
|                  |                |                 |         | asc<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.1.0/HDPDC-GPL-7.1.1.0-centos7-rpm.tar.gz.asc">https://archive.cloudera.com/p/HDPDC-GPL/centos7/7.x/updates/7.1.1.0/HDPDC-GPL-7.1.1.0-centos7-rpm.tar.gz.asc</a> |

| Operating System | Version Number  | Repository Name | Format                        | URL   |
|------------------|-----------------|-----------------|-------------------------------|---|
| Ubuntu 18        | HDP-7.1.1.0-565 | HDPDC           | Version Definition File (VDF) | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.1.0/HDP-7.1.1.0-565.xml">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.1.0/HDP-7.1.1.0-565.xml</a> |
|                  |                 |                 | Base URL                      | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.1.0/HDP%20main">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.1.0/HDP main</a>                     |

| Operating System | Version Number | Repository Name | Format    | URL   |
|------------------|----------------|-----------------|-----------|---|
|                  |                |                 | Repo File | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.1.0/hdpdc.list">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.1.0/hdpdc.list</a>   |
|                  |                |                 | Tarball   | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.1.0/HDPDC-7.1.1.0-ubuntu18-deb.tar.gz">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.1.0/HDPDC-7.1.1.0-ubuntu18-deb.tar.gz</a>         |
|                  |                |                 | md5       | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.1.0/HDPDC-7.1.1.0-ubuntu18-deb.tar.gz.md5">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.1.0/HDPDC-7.1.1.0-ubuntu18-deb.tar.gz.md5</a> |
|                  |                |                 | asc       | <a href="https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.1.0/HDPDC-7.1.1.0-ubuntu18-deb.tar.gz.asc">https://archive.cloudera.com/p/HDPDC/ubuntu18/7.x/updates/7.1.1.0/HDPDC-7.1.1.0-ubuntu18-deb.tar.gz.asc</a> |
|                  |                | HDP-UTILS       | Base URL  | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/</a>   |
|                  |                |                 | Tarball   | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz</a>         |
|                  |                |                 | md5       | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.md5">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.md5</a> |
|                  |                |                 | asc       | <a href="https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.asc">https://archive.cloudera.com/p/HDP-UTILS/1.1.0.22/repos/ubuntu18/HDP-UTILS-1.1.0.22-ubuntu18.tar.gz.asc</a> |
|                  |                | HDPDC-GPL       | URL       | <a href="https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.1.0/hdpdc.gpl.list">https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.1.0/hdpdc.gpl.list</a>                                       |

| Operating System | Version Number | Repository Name | Format  | URL  |
|------------------|----------------|-----------------|---------|--|
|                  |                |                 | Tarball | <a href="https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.1.0/HDPDC-GPL-7.1.1.0-ubuntu18-deb.tar.gz">https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.1.0/HDPDC-GPL-7.1.1.0-ubuntu18-deb.tar.gz</a>                |
|                  |                |                 |         | md5<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.1.0/HDPDC-GPL-7.1.1.0-ubuntu18-deb.tar.gz.md5">https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.1.0/HDPDC-GPL-7.1.1.0-ubuntu18-deb.tar.gz.md5</a> |
|                  |                |                 |         | asc<br><a href="https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.1.0/HDPDC-GPL-7.1.1.0-ubuntu18-deb.tar.gz.asc">https://archive.cloudera.com/p/HDPDC-GPL/ubuntu18/7.x/updates/7.1.1.0/HDPDC-GPL-7.1.1.0-ubuntu18-deb.tar.gz.asc</a> |

### Software download matrix for HDP 2.6.5 to CDP 7.1.x

All the download links related to HDP, Ambari, Cloudera Manager, Cloudera Private Cloud Base, and so on are available here.

| Product  | Download location   | Note  |
|--|---|---|
| AMS pre-upgrade script downloadable                                | <a href="#">Ambari Metrics Pre-upgrade 7.1.1.0</a> <a href="#">Ambari Metrics Pre-upgrade 7.1.2.0</a> <a href="#">Ambari Metrics Pre-upgrade CentOS 7.1.3.0</a> <a href="#">Ambari Metrics Pre-upgrade Ubuntu 7.1.3.0</a> <a href="#">Ambari Metrics Pre-upgrade CentOS 7.1.4.0</a> <a href="#">Ambari Metrics Pre-upgrade Ubuntu 7.1.4.0</a> <a href="#">Ambari Metrics Pre-upgrade SLES 7.1.4.0</a> <a href="#">Ambari Metrics Pre-upgrade CentOS 7.1.6.0</a> <a href="#">Ambari Metrics Pre-upgrade Ubuntu 7.1.6.0</a> <a href="#">Ambari Metrics Pre-upgrade SLES 7.1.6.0</a> <a href="#">Ambari Metrics Pre-upgrade CentOS 7.1.7.0</a> <a href="#">Ambari Metrics Pre-upgrade Ubuntu 7.1.7.0</a> <a href="#">Ambari Metrics Pre-upgrade SLES 7.1.7.0</a> |   |
| Ambari   | <a href="#">Ambari 7.1.x.0</a>  |   |
| HDP  | <a href="#">HDP 7.1.x.0</a>   |   |
| Cloudera Runtime   | <a href="#">Cloudera Runtime</a>  |   |
| AM2CM 2.8.1.0 tool - This tool consists of log4j security updates. | <a href="#">AM2CM 2.8.1.0</a>   | This latest AM2CM tool supports the upgrades from: <ol style="list-style-type: none"> <li>1. HDP 2.6.5 to CDP 7.1.7 SP2</li> <li>2. HDP 2.6.5 to CDP 7.1.8</li> <li>3. HDP 2.6.5 to CDP 7.1.7 SP1</li> <li>4. HDP 2.6.5 to CDP 7.1.7</li> <li>5. HDP 2.6.5 to CDP 7.1.6</li> <li>6. HDP 2.6.5 to CDP 7.1.4</li> </ol> |
| AM2CM Legacy tools   | To access the AM2CM legacy tools, see <a href="#">AM2CM legacy tools download</a>   |   |
| Atlas  | <a href="#">Atlas Migration Exporter</a>  | Atlas Migration Exporter is for upgrading from HDP 2 to CDP Private Cloud Base.   |

| Product                                  | Download location                | Note |
|--|----------------------------------|------|
| Cloudera Manager                         | <a href="#">Cloudera Manager</a> |      |
| Cloudera Manager Packages                | <a href="#">Install packages</a> |      |
| Hive Pre-Upgrade Tool JAR for Compaction | <a href="#">Compaction tool</a>  |      |

### AM2CM legacy tools download

All the download links related to the AM2CM legacy tools are available here.

| Product   | Download location   | Note  |
|---|---|---|
| AM2CM 2.4.3 tool - This tool consists of log4j security updates.                                | <a href="#">AM2CM 2.4.3.0</a>   | This latest AM2CM tool supports the upgrades from: <ol style="list-style-type: none"> <li>1. HDP 2.6.5 to CDP 7.1.4</li> <li>2. HDP 2.6.5 to CDP 7.1.6</li> <li>3. HDP 2.6.5 to CDP 7.1.7</li> <li>4. HDP 2.6.5 to CDP 7.1.8</li> </ol> |
| AM2CM 2.4.2 tool - This tool consists of log4j security updates.                                | <a href="#">AM2CM 2.4.2.0</a>   | This latest AM2CM tool supports the upgrades from: <ol style="list-style-type: none"> <li>1. HDP 2.6.5 to CDP 7.1.4</li> <li>2. HDP 2.6.5 to CDP 7.1.6</li> <li>3. HDP 2.6.5 to CDP 7.1.7</li> <li>4. HDP 2.6.5 to CDP 7.1.8</li> </ol> |
| AM2CM 2.4.1 tool - This tool consists of log4j security updates.                                | <a href="#">AM2CM 2.4.1.1</a> and <a href="#">AM2CM 2.4.1.0</a>   | This latest AM2CM tool supports the upgrades from: <ol style="list-style-type: none"> <li>1. HDP 2.6.5 to CDP 7.1.4</li> <li>2. HDP 2.6.5 to CDP 7.1.6</li> <li>3. HDP 2.6.5 to CDP 7.1.7</li> <li>4. HDP 2.6.5 to CDP 7.1.8</li> </ol> |
| AM2CM 2.4.0 tool - This tool consists of log4j security updates.                                | <a href="#">AM2CM 2.4.0.0</a>   | This latest AM2CM tool supports the upgrades from: <ol style="list-style-type: none"> <li>1. HDP 2.6.5 to CDP 7.1.4</li> <li>2. HDP 2.6.5 to CDP 7.1.6</li> <li>3. HDP 2.6.5 to CDP 7.1.7</li> <li>4. HDP 2.6.5 to CDP 7.1.8</li> </ol> |
| AM2CM 2.x.x tools - The AM2CM 2.1.0, 2.0.4, and 2.0.3 tools consists of log4j security updates. | <a href="#">AM2CM 2.1.0 RHEL7</a> <a href="#">AM2CM 2.1.0 Ubuntu18</a> <a href="#">AM2CM 2.1.0 Sles12</a>       | Supports upgrades from: <ol style="list-style-type: none"> <li>1. HDP 2.6.5 to CDP 7.1.4</li> <li>2. HDP 2.6.5 to CDP 7.1.6</li> <li>3. HDP 2.6.5 to CDP 7.1.7</li> </ol>   |
|   | <a href="#">AM2CM 2.0.4 RHEL7</a> <a href="#">AM2CM 2.0.4 Ubuntu18</a> <a href="#">AM2CM 2.0.4 Sles12</a>       |   |
|   | <a href="#">AM2CM 2.0.3 RHEL7</a> <a href="#">AM2CM 2.0.3 Ubuntu18</a> <a href="#">AM2CM 2.0.3 Sles12</a>       |   |
|   | <a href="#">AM2CM 2.0.2 RHEL7</a> <a href="#">AM2CM 2.0.2 Ubuntu18</a> <a href="#">AM2CM 2.0.2 Sles12</a>       |   |
| AM2CM Legacy tools  | <a href="#">AM2CM 2.0.0.0 RHEL7</a> <a href="#">AM2CM 2.0.0.0 Ubuntu18</a> <a href="#">AM2CM 2.0.0.0 Sles12</a> | To upgrade from HDP 2.6.5 to CDP 7.1.7  |
|   | <a href="#">AM2CM 1.1.1.0</a>   | To upgrade from HDP 2.6.5 to CDP 7.1.6  |
|   | <a href="#">AM2CM 1.1.0.0</a>   | To upgrade from HDP 2.6.5 to CDP 7.1.4  |

## Install software on the hosts

If you have successfully registered the software repositories for the new target version with Ambari, you can now install the software on all hosts in the cluster.

### Procedure

1. Browse to Cluster Admin > Stack and Versions.
2. Click the Versions tab.
3. On a registered target version, click Install Packages and click OK to confirm.  
The install version operation starts. This installs the target version on all hosts in the cluster. You can monitor the progress of the installation by clicking the Installing link. When the installation completes, the Upgrade button replaces the Install Packages button.

## Preparing the services for upgrade

You must complete the pre-upgrade steps for the components and services that you are going to upgrade.

### Backing up Ambari infra data

In Ambari 7.1.x.x, Ambari Infra Solr uses Solr 8. For data stored in the existing Solr 5 implementation used by Ambari 2.6.2.x, the data stored in Ambari Infra must be backed up and the Ambari Infra Solr services must be upgraded.

The next section walks you through steps to both back up the data stored in Ambari Infra and upgrade the Ambari Infra components. The backed up data will be migrated and restored after the HDP upgrade is complete.

#### Back up and upgrade Ambari infra and Ambari Log Search

The Ambari Infra Solr instance is used to index data for Atlas, Ranger, and Log Search. The version of Solr used by Ambari Infra in Ambari 2.6.2.x is Solr 5. The version of Solr used by Ambari Infra in Ambari 7.1.x.x is Solr 8. When moving from Solr 5 to Solr 8 indexed data must be backed up from Solr 5, migrated, and restored into Solr 8 as there are on-disk format changes, and collection-specific schema changes.

The Ambari Infra Solr components must also be upgraded. Scripts are available to do both, and are explained below.

This process will be broken up into four steps:

#### Generate Migration Config

The migration utility requires some basic information about your cluster and this step generates a configuration file that captures that information.

#### Back up Ambari Infra Solr Data

This process backs up all indexed data either to a node-local disk, shared disk (NFS mount), or HDFS filesystem.

#### Remove existing collections and Upgrade Binaries

This step removes the Solr 5 collections, upgrades Ambari Infra to Solr 8, and creates the new collections with the upgraded schema required by the HDP intermediate bits services. This step also upgrades LogSearch binaries if they are installed.

#### Migrate and Restore

This step migrates the backed up data to the new format required by Solr 8 and restores the data into the new collections. This step is completed after the HDP intermediate bits upgrade has been completed in the Post-upgrade Steps section of the upgrade guide

#### Generate migration configuration

The utility used in this process is included in the ambari-infra-solr-client package. This package must be upgraded before the utility can be run. To do this:

### Procedure

1. SSH into a host that has an Infra Solr Instance installed on it. You can locate this host by clicking Hosts on the Ambari Web UI. Click on the Filter icon and type Infra Solr Instance: All to find each host that has an Infra Solr Instance installed on it.

2. Upgrade the `ambari-infra-solr-client` package.

- `yum clean all`
- `yum upgrade ambari-infra-solr-client -y`

For Ubuntu18:

- `apt-get clean all`
- `apt-get install -y ambari-infra-solr-client`

3. You can now proceed to configuring and running the migration tool from the same host. Run the following commands as root, or as a user with sudo access: Export the variable that will hold the full path and filename of the configuration file that is generated in the next step.

```
export CONFIG_INI_LOCATION=ambari_solr_migration.ini
```

4. Run the `migrationConfigGenerator.py` script, located in the `/usr/lib/ambari-infra-solr-client/` directory, with the following parameters:

**--ini-file \$CONFIG\_INI\_LOCATION**

This is the previously exported environmental variable that holds the path and filename of the configuration file that will be generated.

**--host=ambari.hortonworks.local**

This should be the hostname of the Ambari Server.

**--port=8080**

This is the port of the Ambari Server. If the Ambari Server is configured to use HTTPS, use the HTTPS port and add the `-s` parameter to configure HTTPS as the communication protocol.

**--cluster=c1**

This is the name of the cluster that is being managed by Ambari. To find the name of your cluster, look in the upper right and corner of the Ambari Web UI, just to the left of the background operations and alerts.

**--username=admin**

This is the name of a user that is an Ambari Admin.

**--password=admin**

This is the password of the user.

**--backup-base-path=/my/path**

This is the location where the backed up data is stored. Data is backed up to this local directory path on each host that is running an Infra Solr instance in the cluster. So, if you have 3 Infra Solr server instances and you use `--backup-base-path=/home/solr/backup`, this directory is created on all 3 hosts and the data for that host is backed up to this path.

If you are using a shared file system that is mounted on each Infra Solr instance in the cluster, use the `--shared-drive` parameter instead of `--backup-base-path`. The value of this parameter should be the path to the mounted drive that is used for the backup. When this option is chosen, a directory is created in this path for each Ambari Infra Solr instance with the backed up data. For example, if you



had an NFS mount `/export/solr` on each host, you would use `--shared-drive=/exports/solr`. Only use this option if this path exists and is shared amongst all hosts that are running the Ambari Infra Solr.

**--java-home=/usr/jdk64/jdk1.8.0\_112**

This should point to a valid Java 1.8 JDK that is available at the same path on each host in the cluster that is running an Ambari Infra Solr instance.

**If the Ranger Audit collection is being stored in HDFS, add the parameter, --ranger-hdfs-base-path.**

The value of this parameter should be set to the path in HDFS where the Solr collection for the Ranger Audit data has been configured to store its data.

Example: `--ranger-hdfs-base-path=/user/infra-solr`

Example Invocations:

If using HTTPS for the Ambari Server:

```
/usr/bin/python /usr/lib/ambari-infra-solr-client/migrationConfigGenerator.py \
--ini-file $CONFIG_INI_LOCATION \
--host=c7401.ambari.apache.org \
--port=8443 -s \
--cluster=c11 \
--username=admin \
--password=admin \
--backup-base-path=/home/solr/backup \
--java-home=/usr/jdk64/jdk1.8.0_112
```

If using HTTP for the Ambari Server:

```
/usr/bin/python /usr/lib/ambari-infra-solr-client/migrationConfigGenerator.py \
--ini-file $CONFIG_INI_LOCATION \
--host=c7401.ambari.apache.org \
--port=8080 \
--cluster=c11 \
--username=admin \
--password=admin \
--backup-base-path=/home/solr/backup \
--java-home=/usr/jdk64/jdk1.8.0_112
```

Ensure the script `migrationConfigGenerator.py` generates cleanly and there are no yellow warning texts visible. If so, review the yellow warnings.

### Back up Ambari Infra Solr data

After the configuration file is generated, you must review the ini file created by the process.

### About this task

There is a configuration section for each collection that was detected. If you do not want to backup a specific collection you can set `enabled = false` and the collection will not be backed up. Ensure that `enabled = true` is set for all of the collections you do wish to back up. By default, only the Atlas and Ranger collections will be backed up and Log Search will not be backed up..

### Procedure

- To execute the backup, run the following command from the same host on which you generated the configuration file:

```
# /usr/lib/ambari-infra-solr-client/ambariSolrMigration.sh \
--ini-file $CONFIG_INI_LOCATION \
```

```
--mode backup | tee backup_output.txt
```

During this process, the script generates Ambari tasks that are visible in the Background Operations dialog in the Ambari Server. Once the process has completed, retain the output of the script for your records. This output is helpful when debugging any issues that may occur during the migration process, and the output contains information regarding the number of documents and size of each backed up collection.

### Remove Existing Collections and Upgrade Binaries

Once the data has been backed up, the old collections need to be deleted, and the Ambari Infra Solr, and Log Search (if installed) components need to be upgraded.

#### Procedure

- To do all of that, run the following script:

```
/usr/lib/ambari-infra-solr-client/ambariSolrMigration.sh \  
--ini-file $CONFIG_INI_LOCATION \  
--mode delete | tee delete_output.txt
```

During this process, the script generates Ambari tasks that are visible in the Background Operations dialog in the Ambari Server.

After the process has completed, retain the output of the script for your records. This output is helpful when debugging any issues that may occur during the migration process.

### Upgrading Ambari Infra

If you have Ambari Solr installed, you must upgrade Ambari Infra after upgrading Ambari.

#### Procedure

1. Make sure Ambari Infra services are stopped. From Ambari Web, browse to Services > Ambari Infra and select Stop from the Service Actions menu.
2. On every host in your cluster with an Infra Solr Client installed, run the following commands:

For RHEL/CentOS/Oracle Linux:

```
yum clean all  
yum upgrade ambari-infra-solr-client
```

For SLES:

```
zypper clean  
zypper up ambari-infra-solr-client
```

For Ubuntu/Debian:

```
apt-get clean all  
apt-get update  
apt-get install ambari-infra-solr-client
```

3. Execute the following command on all hosts running an Ambari Infra Solr Instance:

For RHEL/CentOS/Oracle Linux:

```
yum upgrade ambari-infra-solr
```

For SLES:

```
zypper up ambari-infra-solr
```

For Ubuntu/Debian:

```
apt-get install ambari-infra-solr
```

4. Start the Ambari Infra services.

From Ambari Web, browse to Services > Ambari Infra select Service Actions then choose Start.

## Overview of the Migration of the Atlas and Infra Solr Data

The illustration helps you to understand the overview of the upgrade of the the Atlas and Infra Solr services' data from HDP to CDP Private Cloud Base.

### Preparing Atlas for upgrade

The Atlas migration exporter utility is used for exporting data from Atlas from HDP 2.6.5.x. You must import this data into the CDP Private Cloud Base version of Atlas.

To estimate the duration for the export operation, use one of the following methods to determine the number of entities in Atlas: The time the tool takes to export the Atlas metadata depends on the number of entities and your cluster configuration.

- Within Solr, check the number of entities in the vertex\_index collection.

- Run the following Atlas metrics REST API query, with your values for user, password, and Atlas port:

```
curl -g -X GET -u admin:admin -H "Content-Type: application/json" -H"Cache-Control: no-cache" "http://<atlas_server>:21000/api/atlas/admin/metrics"
```



**Note:** Estimated duration for export: Two million entities per hour, for a node with a 8 GB RAM quad-core processor with both the Atlas and Solr servers on the same node.

If your environment is kerberos enabled, perform the following steps. and proceed to exporting Atlas data.



**Attention:** If your environment is not kerberos enabled, then skip the steps below and proceed to exporting Atlas data.

1. Switch identity to the Atlas user. `$ su - atlas`
2. List the Atlas keytab to find the Atlas Kerberos Service Principal value.

```
$ klist -kt /etc/security/keytabs/atlas.service.keytab
Keytab name: FILE:/etc/security/keytabs/atlas.service.keytab
KVNO Timestamp Principal
-----
-----
1111/28/2018 06:40:24 atlas/finance.acme.trade@ACME.TRADE
1 11/28/2018 06:40:24 atlas/finance.acme.trade@ACME.TRADE
```

3. Initialize the Atlas keytab and the Atlas Kerberos Service Principal value.

```
$ kinit -kt /etc/security/keytabs/atlas.service.keytab atlas/finance.acme.trade@ACME.TRADE
```

4. Verify you have a valid ticket.

```
$ klist
Ticket cache: FILE:/tmp/krb5cc_1005
Default principal: atlas/finance.acme.trade@ACME.TRADE
Valid starting Expires Service principal
11/29/2020 22:52:56 11/30/2020 22:52:56 krbtgt/ACME.TRADE@ACME.TRADE
renew until 12/06/2020 22:52:56
```

Export Atlas data:

The Atlas migration exporter utility is used for migrating Atlas from HDP 2.6.5.x to HDP 7.1.x:

1. In the Ambari Web UI, click Atlas.
2. Select Actions>Stop.
3. Ensure that HBase is running. If it is not running, in the Ambari Web UI, click HBase
4. Select Actions>Start.
5. SSH to the host where your Atlas Metadata Server is installed.
6. Download from the location: [Atlas migration exporter](#).
7. Create this directory `mkdir /usr/hdp/2.6.<current version>/atlas/tools/migration-exporter/`
8. Copy the contents to `/usr/hdp/2.6.<current version>/atlas/tools/migration-exporter/`
9. Modify the permissions using `chown -R atlas:atlas <directory above>`

10. Execute the tool from the location above. `atlas_migration.py -d <output directory>`. For example, `/usr/hdp/2.6.<current version>/atlas/tools/migration-exporter/atlas_migration.py -d /var/lib/atlas-data`

For example, this command exports the metadata to an output directory called `atlas_metadata`. `python /usr/hdp/current/atlas/tools/migration-exporter/atlas_migration_export.py -d ./atlas_metadata`



**Note:** While running the Atlas export migration tool, the tool prevents the use of Atlas, and blocks all REST APIs and Atlas hook notification processing. The time the tool takes to export the Atlas metadata depends on the number of entities and your cluster configuration.

11. To display the export status, use this command: `tail -f /var/log/atlas/atlas-migration-exporter.log`.



**Note:** The export task creates a file named `atlas-migration-data.json` and `atlas-migration-typesdef.json` in the specified output directory specified, which contains the exported Atlas entity data. Preserve this output for later steps to populate Atlas.

### Place Atlas in migration mode

Before starting the upgrade from HDP 2.6.5.x to HDP intermediate bits, within Atlas you must set the `atlas.migration.data.filename` property to point to a non-existent directory.

It is mandatory to ensure that the pre-upgrade checks within Ambari are completed. After you upgrade to HDP intermediate bits, the migration is prevented from starting. This property will be set to the correct path. Configure the `atlas.migration.data.filename` property.

1. In Ambari, select Services > Atlas > Configs > Advanced > Custom application-properties.
2. Click Add Property, and add the `atlas.migration.data.filename` property.

Example: `atlas.migration.data.filename=/tmp/non-existent-path`

### Ranger Service connection with Oracle database

If you have configured Ranger with Oracle Database as its backend, you must add the property with Key and Value.

#### Ranger connection with Oracle DB

If you have configured Ranger with Oracle Database as its backend, you must add the property `conf/ranger-admin-site.xml` with the key and value provided below for the Custom `ranger-admin-site` section in Ambari: Key: `ranger.jpa.jdbc.preferredtestquery` value: `select 1 from dual;`

### Ranger admin password

During the HDP upgrade, Ranger admin fails to apply the Java patch to change all admin passwords if the `hadoop.security.credstore.java-keystore-provider.password-file` property is set in `core-site`.

#### About this task

Before you upgrade from HDP 2.6.5 to HDP intermediate bits, you must manually perform the following steps:

#### Procedure

1. SSH to the ranger admin host
2. `cd /etc/ranger/admin/conf/`
3. `vi ranger-admin-env-credstore.sh`
4. add `"export HADOOP_CREDSTORE_PASSWORD=none"` in the `"ranger-admin-env-credstore.sh"` file
5. `chown ranger:ranger ranger-admin-env-credstore.sh`
6. `chmod 755 ranger-admin-env-credstore.sh`

## Preparing HBase for upgrade

Before upgrading to Cloudera Private Cloud Base, transition your HFiles to a supported encoding type and validate HFiles to ensure that you do not have any HFiles that use unsupported encoding or are in a bad state.



**Important:** Ensure that you complete all the pre-upgrade steps if you have Apache HBase installed in your existing HDP cluster.

### Preparing HBase for upgrade

You must perform the pre-upgrade tasks after installing the HDP intermediate bits on the hosts, but before you upgrade your cluster to HDP intermediate bits. Run these commands on the cluster that has the HDP intermediate bits. The HBase validation tool is packaged within HDP intermediate bits, and you register the software repositories for the new target version with Ambari and then install the software on all hosts in the cluster. For more information to get the intermediate bits, see [Register software repositories](#).

Ensure that you check for any corrupted HFiles and rectify them before the upgrade using the HFile content validator (validate-hfile) tool.

When upgrading from an HDP 2.6.x cluster, you should begin by making a copy of the HBase client configuration directory for use with these upgrade checks. In this configuration copy, you should delete the metrics collection configuration file `hadoop-metrics2-hbase.properties` because it is incompatible with the newer utilities. If this file is left in place then the pre-upgrade HBase commands will not properly check the underlying HFiles for issues. Before you run the pre-upgrade checks, ensure that you point to the correct configuration directory. For example, you can use the `export` command or the `hbase --config` CLI option to point to the configuration directory.

```
export HBASE_CONF_DIR=[**PATH TO THE CONFIGURATION DIRECTORY**]
```

This is an example `export HBASE_CONF_DIR` command that you can use to point to the correct configuration directory before you run the validation checks.

```
export HBASE_CONF_DIR=/home/dirname/hbase-conf
```

Should the pre-upgrade check fail due to an old incompatible configuration directory, the command will log the following exception:

```
java.lang.AbstractMethodError: org.apache.hadoop.metrics2.sink.timeline.HadoopTimelineMetricsSink.init
(Lorg/apache/commons/configuration2/SubsetConfiguration;)
```

### What's next:

- Remove PREFIX\_TREE data block encoding: `hbase pre-upgrade validate-dbe`
- Validate HFiles: `hbase pre-upgrade validate-hfile`

### Remove PREFIX\_TREE data block encoding

Before upgrading to Cloudera Private Cloud Base, you must ensure that you have transitioned all the data to a supported encoding type.

### Before you begin

- If your cluster is Kerberized, ensure that you run the `kinit` command as a `hbase` user before running the pre-upgrade commands.
- Ensure you have valid Kerberos credentials. You can list the Kerberos credentials using the `klist` command, and you can obtain credentials using the `kinit` command.



**Important:** Ensure that you complete all the pre-upgrade steps listed in *Preparing HBase for upgrade* if you have Apache HBase installed in your existing HDP cluster.

## About this task

Use the `hbase pre-upgrade validate-dbe` command to identify HFiles using the `PREFIX_TREE` data block encoding.



**Important:** In HDP, you could use the `PREFIX_TREE` data block encoding as an unsupported option. But, `PREFIX_TREE` is completely removed in Cloudera Private Cloud Base. Before upgrading, ensure that you have transitioned your HFiles to a supported data encoding type.

## Procedure

1. Run the `hbase pre-upgrade validate-dbe` command on the cluster that has the HDP intermediate bits.

For example, if you are using the HDP 7.1.2.0-96 intermediate bits, you must run the following command:

```
/usr/hdp/7.1.2.0-96/hbase/bin/hbase pre-upgrade validate-dbe
```

The commands check whether your table or snapshots use the `PREFIX_TREE` data block encoding. This command does not take much time to run because it validates only the table-level descriptors.

If the HFiles do not use `PREFIX_TREE` data block encoding, the following message is displayed:

```
The used Data Block Encodings are compatible with HBase 2.0.
```

If you see this message, your data block encodings are compatible, and you do not have to do any more steps.

If an HFile uses the `PREFIX_TREE` data block encoding, a similar error message to the following is displayed:

```
2018-07-13 09:58:32,028 WARN [main] tool.DataBlockEncodingValidator: In compatible DataBlockEncoding for table: t, cf: f, encoding: PREFIX_TREE
```

If you see this error message, continue to Step 2 and fix all your `PREFIX_TREE` encoded tables.

2. Fix your `PREFIX_TREE` encoded tables using the old installation.

You can change the Data Block Encoding type to `PREFIX`, `DIFF`, or `FAST_DIFF` in your source cluster.

For example, if your validation output reported column family `f` of table `t` is invalid. Its Data Block Encoding type is changed to `FAST_DIFF`:

```
hbase> alter 't', { NAME => 'f', DATA_BLOCK_ENCODING => 'FAST_DIFF' }
```

- a) Run the following command on the cluster that has the HDP intermediate bits.

```
hbase pre-upgrade validate-hfile
```



**Note:** This command checks every HFile (including snapshots) to ensure none of the HFile have `PREFIX_TREE` data block encoding. It opens every HFile, so this operation can take a long time.

If there is no HFile with `PREFIX_TREE` encoding, a confirmation message similar to the following is displayed:

```
Checked N HFiles, none of them are corrupted.
There are no incompatible HFiles.
```

If you have an HFile with `PREFIX_TREE` encoding, an error message similar to the following is displayed:

```
INFO [main] tool.HFileContentValidator: Corrupted file: hdfs://example.com:8020/hbase/data/default/t/72ea7f7d625ee30f959897d1a3e2c350/prefix/7e6b3d73263c4851bf2b8590a9b3791e
```

```
2018-06-05 16:20:47,383 INFO [main] tool.HFileContentValidator: Corrupted file: hdfs://example.com:8020/hbase/archive/data/default/t/56be41796340b757eb7fff1eb5e2a905/f/29c641ae91c34fc3bee881f45436b6d1
```

If the error message is displayed, go to the next step. If you don't see an error, skip to *Validate HFile*.

- b) Get the table name from the output. Our example output showed the `/hbase/data/default/t/...` file path, which means that the HFile belongs to the `t` table which is in the default namespace.
- c) Rewrite the HFiles by running major compaction in the old installation using this command.

```
hbase > major_compact 't'
```

When the major compaction is complete, the original invalid HFile is automatically removed.

- d) Run the `hbase pre-upgrade validate-hfile` command again to confirm if all the HFiles are rewritten in the new encoding format.

### What to do next

Validate HFile

### Validate HFile

Before upgrading to Cloudera Private Cloud Base, ensure that your HFiles are compatible with the version Apache HBase in CDP.

### Before you begin

If your cluster is Kerberized, ensure that you run the `kinit` command as a `hbase` user before running the pre-upgrade commands.



**Important:** Ensure that you complete all the pre-upgrade steps listed in *Preparing HBase for upgrade* if you have Apache HBase installed in your existing HDP cluster.

### About this task

After converting all the `PREFIX_TREE` encoded HFiles to a supported format, there may be HFiles that are not compatible with HBase in CDP. Run the `validate-hfile` command again on the cluster that has the HDP intermediate bits. Also check for bad state of HFiles using the `hbck` command.

Use the following pre-upgrade commands to validate HFiles:

- ```
hbase pre-upgrade validate-hfile
```
- ```
hbase hbck
```

The `validate-hfile` tool lists all the corrupt HFiles with details. The `hbck`

### Before you begin

Use the CDP parcel to run the pre-upgrade commands. Cloudera recommends that you run them on an HMaster host.



## Procedure

1. Run the `hbase pre-upgrade validate-hfile` command on the cluster that has the HDP intermediate bits.

This command checks every HFile (including snapshots) to ensure that the HFiles are not corrupted and have the compatible block encoding. It opens every HFile, and this operation can take a long time based on the size and number of HFiles in your cluster.

For example, if you are using the HDP 7.1.2.0-96 intermediate bits, you must run the following command:

```
/usr/hdp/7.1.2.0-96/hbase/bin/hbase pre-upgrade validate-hfile
```

If there are no incompatible HFiles, the following message is displayed:

```
Checked 3 HFiles, none of them are corrupted.
There are no incompatible HFiles.
```

If you have an incompatible HFile, a similar error message to this is displayed:

```
2018-06-05 16:20:47,322 INFO [main] tool.HFileContentValidator: Corrupted
file: hdfs://example.com:8020/hbase/data/default/t/72ea7f7d625ee30f959897
d1a3e2c350/prefix/7e6b3d73263c4851bf2b8590a9b3791e
2018-06-05 16:20:47,383 INFO [main] tool.HFileContentValidator: Corrupte
d file: hdfs://example.com:8020/hbase/archive/data/default/t/56be4179634
0b757eb7fff1eb5e2a905/f/29c641ae91c34fc3bee881f45436b6d1
```

If you get an error message, continue with Step 2 otherwise skip to Step 5.

2. Identify what kind of HFiles were reported in the error message.

The report can contain two different kinds of HFiles, they differ in their path:

- If an HFile is in `/hbase/data` then it belongs to a table.
- If an HFile is located under `/hbase/archive/data` then it belongs to a snapshot.

3. Fix the HFiles that belong to a table.

- a) Get the table name from the output.

Our example output showed the `/hbase/data/default/t/...` path, which means that the HFile belongs to the `t` table which is in the default namespace.

- b) Rewrite the HFiles by issuing a major compaction. Use the old installation.

```
shell> major_compact 't'
```

When compaction is finished, the invalid HFile disappears.

4. Fix the HFiles that belong to a snapshot that is needed. Use the old installation.

- a) Find the snapshot which refers the invalid HFile. You can do this using a shell script and the following example code. In our example output, the HFile is `29c641ae91c34fc3bee881f45436b6d1`:

```
#!/bin/bash
# Find the snapshot which referes to the invalid HFile
for snapshot in $(hbase snapshotinfo -list-snapshots 2> /dev/null | cut
-f 1 -d \|);
do
echo "checking snapshot named '${snapshot}'"
hbase snapshotinfo -snapshot "${snapshot}" -files 2> /dev/null | grep
29c641ae91c34fc3bee881f45436b6d1
done
```

The following output means that the invalid file belongs to the `t_snap` snapshot:

```
checking snapshot names 't__snap'
1.0 K t/56be41796340b757eb7fff1eb5e2a905/f/29c641ae91c34fc3bee881f454
36b6d1 (archive)
```

- b) Convert snapshot to another HFile format if data encoding is the issue using the following command:

```
# creating a new namespace for the cleanup process
hbase> create_namespace 'pre_upgrade_cleanup'
# creating a new snapshot
hbase> clone_snapshot 't_snap', 'pre_upgrade_cleanup:t'
hbase> alter 'pre_upgrade_cleanup:t', { NAME => 'f', DATA_BLOCK_ENCODING => 'FAST_DIFF' }
hbase> major_compact 'pre_upgrade_cleanup:t'
```



**Important:** Confirm if the major compaction is complete in the RegionServer Web user interface or the RegionServer logs before you run the following commands to remove invalid snapshots. If the major compaction process is not completed when you delete the snapshot, your files may get corrupted.

```
# removing the invalid snapshot
hbase> delete_snapshot 't_snap'
# creating a new snapshot
hbase> snapshot 'pre_upgrade_cleanup:t', 't_snap'
# removing temporary table
hbase> disable 'pre_upgrade_cleanup:t'
hbase> drop 'pre_upgrade_cleanup:t'
hbase> drop_namespace 'pre_upgrade_cleanup'
```

5. Run the hbck command on to identify HFiles in a bad state and remedy those HFiles.

## Preparing HDP Search for upgrade

Cloudera provides the solr-upgrade.sh script to transition HDP Search 2 or 3 configurations to make them compatible with Solr 8 that is shipped with Cloudera Runtime 7.1.1 or higher. To run this script, you need to download configuration metadata from your HDP Search instance and copy it to an interim CDP cluster with Solr service. You can then use that same cluster to validate and test the migrated configuration metadata.



### Note:

Do not perform the upgrade tasks on a production cluster, during an actual HDP to CDP upgrade, because the process is time consuming and may require fixing incompatibilities manually.

The migration script cannot upgrade the Lucene index files. After upgrading, you must reindex your collections. For more information, see [Reindexing in Solr](#) in the Apache Solr wiki.

## HDP Search to Cloudera Search transition process

1. Deploy an interim CDP cluster with Cloudera Runtime 7.1.1 or higher. Start a Solr service.
2. Download the HDP Search configuration and copy it to the interim CDP cluster.
3. Transition the HDP Search configuration using the interim cluster.
4. Validate the transitioned configuration, using the interim cluster.
5. Test the transitioned configuration, using the interim cluster.
6. Upgrade HDP to Cloudera Runtime 7.1.x.
7. Perform *Search post-HDP-upgrade tasks*.

## Warnings and prerequisites

Preparing HDP Search for a CDP Private Cloud Base upgrade is a complex task.



**Warning:** The procedure documented here is not applicable to transition infra-Solr deployments from HDP to CDP.



**Warning:** Due to changes in the underlying Apache Lucene file format between the HDP Search 2 or 3 and the Solr shipped with Cloudera Runtime 7.1.1 or higher, it is not possible to directly upgrade your existing indexes. The following procedures update only your Solr configuration and metadata to be compatible with the new Solr version. After upgrading to Cloudera Runtime 7.1.1 or higher, you must reindex your collections.

Make sure that the source data that was used to index your collections is still available before upgrading. Plan downtime for any applications or services that use your Search deployment.



**Warning:**

The `solr-upgrade.sh` script shipped with Cloudera Manager 7.3.1 or earlier downloads `aliases.json` from ZooKeeper during the pre-upgrade transition, but it fails to upload the file to the upgraded cluster. You need to recreate aliases manually on the upgraded cluster. For more information on aliases, see [Collection aliasing](#).

If you run Cloudera Manager 7.4.2 or higher, you can disregard this warning.



**Warning:** Although the configuration transformation script addresses known incompatibilities, there might be incompatible configurations that are not detected by the script. Do not upgrade your production Cloudera Search services to Cloudera Runtime 7.1.1 or higher until you have tested the migrated configuration with Cloudera Runtime 7.1.1 or higher on a testing or development cluster.

### Prerequisites

- Obtain the `solr-download-metadata.sh` script by creating a support ticket at the [support site](#).
- Provision an interim Cloudera Runtime 7 cluster with a Solr service before you start the upgrade process. You need it to perform some of the upgrade transformation steps and you can also use it for upgrade testing.
- Stop making changes to your HDP Search environment. Make sure that no configuration changes are made to HDP Search for the duration of the transition and upgrade. This includes adding or removing Solr Server hosts, moving Solr Server roles between hosts, changing hostnames, and so on.
- Plan for an outage for any applications or services that use your Search deployment until you have completed reindexing after the upgrade.
- If you use Kerberos, create a `jaas.conf` file for the Search service superuser (`solr` by default).
- Do not create, delete, or modify any collections for the duration of the transition and upgrade.
- Make a backup of your HDP Search deployment.

### Download Solr configuration from HDP Search ZooKeeper

The `solr-download-metadata.sh` script downloads metadata from HDP Search ZooKeeper. You can then copy this information into a CDP cluster and use it as input for the CDP Private Cloud Base upgrade process.

### Before you begin

The migration script uses `jq` for JSON file processing, this tool has to be installed:

For example, `yum install jq`

### Procedure

1. Copy the `solr-download-metadata.sh` script to any of the nodes in the cluster where HDP Search is deployed.

```
#!/bin/bash
exec 3>/dev/null

set -e

ZKCLI_HDP_LUCIDWORKS=/opt/lucidworks-hdpsearch/solr/server/scripts/cloud-
scripts/zkcli.sh
ZKCLI_HDP_CLOUDERA=/usr/cloudera-hdp-solr/current/cloudera-hdp-solr/solr/
server/scripts/cloud-scripts/zkcli.sh
ZKCLI_CDH=/opt/cloudera/parcels/CDH/lib/solr/bin/zkcli.sh
```

```

error() {
    >&2 echo "$1"
    exit 1
}
remove_trailing_slash() {
    echo "$1" | sed 's@/*$@@'
}

check_jq() {
    if ! command -v jq &> /dev/null; then
        echo "Command jq could not be found"
        exit
    fi
}
check_java() {
    if ! command -v java &> /dev/null; then
        echo "Command java could not be found"
        exit
    fi
}

check_zkcli() {
    if [ ! -f ${ZKCLI_HDP_LUCIDWORKS} ]; then
        if [ ! -f ${ZKCLI_HDP_CLOUDERA} ]; then
            if [ ! -f ${ZKCLI_CDH} ]; then
                error "Cannot find zkcli.sh"
            else
                zkcli=${ZKCLI_CDH}
            fi
        else
            zkcli=${ZKCLI_HDP_CLOUDERA}
        fi
    else
        zkcli=${ZKCLI_HDP_LUCIDWORKS}
    fi
    echo "Found zkcli.sh at ${zkcli}"
}

run_zk_cli() {
    : ${SOLR_ZK_ENSEMBLE:? "Please configure Solr ZooKeeper ensemble using
the SOLR_ZK_ENSEMBLE env variable"}
    ${zkcli} -zkhost "$SOLR_ZK_ENSEMBLE" "$@" 2>&3
}

get_collections() {
    if ! jq -r 'to_entries[] | .key' < "$1"; then
        error "Cannot find collections in file $1. Please check the contents
of the file."
    fi
}

get_configname() {
    jq -r '["'$2'"].configName' < "$1"
}

download_zk_metadata() {
# echo "Cleaning up $1"
# rm -rf "${1:?}"/

    mkdir -p "$1/collections"
    mkdir -p "$1/configs"

    echo "Copying clusterstate.json"
# run_zk_cli -cmd get /clusterstate.json > "$1"/clusterstate.json;

```

```

: ${SOLR_ADMIN_URI:? "Please configure Solr URI using the SOLR_ADMIN_URI
env variable"}
SOLR_ADMIN_URI=$(remove_trailing_slash "$SOLR_ADMIN_URI")
curl --retry 5 -s -L -k --negotiate -u : "${SOLR_ADMIN_URI}/admin/col
lections?action=CLUSTERSTATUS&wt=json" 2> /dev/null | jq -r '.cluster.co
llections' > "$1"/clusterstate.json
echo "Copying clusterprops.json"
if ! run_zk_cli -cmd get /clusterprops.json > "$1"/clusterprops.json;
then
echo "Missing/empty clusterprops.json file in ZooKeeper, creating e
mpty json"
echo "{}" > "$1"/clusterprops.json
fi

echo "Copying solr.xml"
if ! run_zk_cli -cmd get /solr.xml > "$1"/solr.xml; then
echo "Missing/empty solr.xml in ZK, copying from local dir"
if ! cp /etc/solr/home/solr.xml "$1"/solr.xml; then
if ! cp /etc/solr/data_dir/solr.xml "$1"/solr.xml; then
if ! cp /opt/solr/data/solr.xml "$1"/solr.xml; then
if ! cp /opt/lucidworks/solr.xml "$1"/solr.xml; then
: ${SOLR_XML_PATH:? "Cannot download solr.xml from ZooKeeper a
nd cannot find it locally. Please specify its location using the SOLR_XM
L_PATH variable"}
if ! cp "${SOLR_XML_PATH}" "$1"/solr.xml; then

error "Cannot find solr.xml. Exiting."
fi
fi
fi
fi
fi
fi
fi

echo "Copying aliases.json"
if ! run_zk_cli -cmd get /aliases.json > "$1"/aliases.json ; then
echo "Unable to copy aliases.json. Please check if it contains any dat
a ?"
echo "Continuing with the download..."
fi

collections=$(get_collections "$1"/clusterstate.json)
for c in ${collections}; do
echo "Downloading configuration for collection ${c}"
run_zk_cli -cmd get /collections/"$c" > "$1/collections/${c}_config
.json"
coll_conf=$(get_configname "$1/clusterstate.json" "$c")
echo "Downloading config named ${coll_conf} for collection ${c}"
run_zk_cli -cmd downconfig -confdir "$1/configs/${coll_conf}/conf" -
confname "${coll_conf}"
done

echo "Successfully downloaded Solr metadata from ZooKeeper"
}
[[ -z "$1" ]] && error "Please specify output directory"

check_jq
check_java
check_zkcli
download_zk_metadata "$1"

```

2. Set the JAVA\_HOME environment variable to the JDK location. For example:

```
export JAVA_HOME=/usr/jdk64/jdk1.8.0_112
```

3. Add java to PATH. For example:

```
export PATH=$PATH:$JAVA_HOME/bin
```

4. On a host running Solr server, set the SOLR\_ZK\_ENSEMBLE environment variable:

```
export SOLR_ZK_ENSEMBLE=[***HOSTNAME***]:2181/solr
```

Replace [\*\*\*HOSTNAME\*\*\*] with a value that is valid in your environment.

5. On a host running Solr server, set the SOLR\_ADMIN\_URI environment variable:

```
export SOLR_ADMIN_URI=[***HOSTNAME***]:8983/solr
```

Replace [\*\*\*HOSTNAME\*\*\*] with an actual value that is valid in your environment. Make sure that you include the protocol (HTTP or HTTPS) and do not include a trailing slash (/).

For example:

```
export SOLR_ADMIN_URI=http://example.com:8983/solr
```

6. On a secure cluster you may need to set additional zkcli flags, for instance:

```
export ZKCLI_JVM_FLAGS="-Djava.security.auth.login.config=[***/PATH/TO/JAAS.CONF***]"
```

Replace [\*\*\*/PATH/TO/JAAS.CONF\*\*\*] with path to a valid jaas.conf file.

7. If you have enabled Kerberos, run the kinit command with the Solr service superuser. For example:

```
kinit solr@[***EXAMPLE.COM***]
```

Replace [\*\*\*EXAMPLE.COM\*\*\*] with your Kerberos realm name.

8. Download the current (HDP Search 2 or 3) Solr configuration from ZooKeeper:

- ```
solr-download-metadata.sh [***DIRECTORY***]
```

Replace [\*\*\*DIRECTORY\*\*\*] with the path to the directory where you want to download the Solr configuration.

- If you have enabled Kerberos and configured ZooKeeper Access Control Lists (ACLs), specify your JAAS configuration file by adding the --jaas parameter to the command.

```
solr-download-metadata.sh --jaas [***/PATH/TO/JAAS.CONF***] [***DIRECTORY***]
```

Replace [\*\*\*/PATH/TO/JAAS.CONF\*\*\*] with the path to a valid jaas.conf file.

Replace [\*\*\*DIRECTORY\*\*\*] with the path to the directory where you want to download the Solr configuration.

In case of a successful run, the last line printed out by the script is:

Successfully downloaded SOLR metadata from Zookeeper



**Note:** Make a backup copy of the downloaded configuration.

9. Copy the migrated data to \$HOME/cr7-solr-migration in the interim Solr instance.

10. Initialize a directory for the migrated configuration as a copy of the current configuration:

```
cp -r $HOME/cr7-solr-migration $HOME/cr7-migrated-solr-config
```

### Transition Solr configuration

Depending on the Cloudera Manager version, select the transition process that is applicable to your case.

#### Transition Solr configuration using Cloudera Manager versions 7.1.1 to 7.2.4

The migration script transforms configuration metadata before the actual upgrade and checks its validity for the target Solr version. In case it identifies incompatibilities it cannot resolve, the script stops, letting you fix the input file with the incompatibility. Afterwards, you can rerun the script to continue with the transition process.



**Warning:** Although the configuration transformation script addresses known incompatibilities, there might be incompatible configurations that are not detected by the script. Do not upgrade your production Cloudera Search services to Cloudera Runtime 7.1.1 or higher until you have tested the migrated configuration with Cloudera Runtime 7.1.1 or higher on a testing or development cluster.

1. Run the migration script:

- a. 

```
/opt/cloudera/cm/solr-upgrade/solr-upgrade.sh config-upgrade -t \  
solrxml -c $HOME/cr7-solr-migration/solr.xml -u \  
/opt/cloudera/cm/solr-upgrade/validators/solr_4_to_8_processors.xml \  
-d /tmp
```

If Kerberos is enabled, specify your JAAS configuration file by appending `--jaas /PATH/TO/SOLR-JAAS.CONF` to the command.

- b. If the script reports any incompatibilities, fix them in the working directory (`$HOME/cr7-solr-migration/solr.xml` in this example) and then rerun the script. Each time you run the script, the files in the output directory (`/tmp` in this example) are overwritten. Repeat until there are no incompatibilities and the `solr.xml` migration is successful. For example:

```
Validating solrxml...  
No configuration errors found...  
No configuration warnings found...  
  
Following incompatibilities will be fixed by auto-transformations  
(using --upgrade command):  
  * System property used to define SOLR server port has changed from  
    solr.port to jetty.port  
  
Solr solrxml validation is successful. Please review  
/tmp/solrxml_validation.html for more details.  
  
Applying auto transformations...  
  
The upgraded configuration file is available at /tmp/solr.xml
```

2. Copy the migrated `solr.xml` file to the migrated configuration directory:

```
cp /tmp/solr.xml $HOME/cr7-migrated-solr-config
```

3. For each collection configuration set in `$HOME/cr7-solr-migration/configs/`, migrate the configuration and schema. Each time you run the script, the output file is overwritten. Do not proceed to the next collection until the migration is successful and you have copied the migrated file to its final destination.

- a. Run the migration script for `solrconfig.xml`. For example, for a configuration set named `TWEETS_CONFIG`:

```
/opt/cloudera/cm/solr-upgrade/solr-upgrade.sh config-upgrade -t \  
solrconfig -c $HOME/cr7-solr-migration/configs/tweets_config/conf/solr  
config.xml -u \  
/opt/cloudera/cm/solr-upgrade/validators/solr_4_to_8_processors.xml \  
-d /tmp
```

- b. If the script reports any incompatibilities, fix them in the working directory (`$HOME/cr7-solr-migration/configs/tweets_config/conf/solrconfig.xml` in this example) and then rerun the script. Repeat until there are no incompatibilities and the `solrconfig.xml` migration is successful. You should see a message similar to the following:

```
Solr solrconfig validation is successful. Please review
/tmp/solrconfig_validation.html for more details.

Applying auto transformations...

The upgraded configuration file is available at /tmp/solrconfig.xml
```

- c. Copy the migrated `solrconfig.xml` file to the collection configuration directory in the migrated directory. For example:

```
cp /tmp/solrconfig.xml \
$HOME/cr7-migrated-solr-config/configs/tweets_config/conf/
```

- d. Run the migration script for `schema.xml` (or `managed-schema`). For example, for a configuration set named `TWEETS_CONFIG`:

```
/opt/cloudera/cm/solr-upgrade/solr-upgrade.sh config-upgrade \
-t schema \
-c $HOME/cr7-solr-migration/configs/tweets_config/conf/schema.xml \
-u /opt/cloudera/cm/solr-upgrade/validators/solr_4_to_8_processors.xml \
-d /tmp
```

- e. If the script reports any incompatibilities, fix them in the working directory (`$HOME/cr7-solr-migration/configs/tweets_config/conf/schema.xml` in this example) and then rerun the script. Repeat until there are no incompatibilities and the `solrconfig.xml` migrations are all successful.
- f. Copy the migrated `schema.xml` file to the collection configuration directory in the migrated directory. For example:

```
cp /tmp/schema.xml $HOME/cr7-migrated-solr-config/configs/tweets_config/
conf/
```

- g. Repeat for all configuration sets in `$HOME/cr7-solr-migration/configs/`.

### Transition the configuration using Cloudera Manager versions 7.3.1 or higher

The migration script transforms configuration metadata before the actual upgrade and checks its validity for the target Solr version. In case it identifies incompatibilities it cannot resolve, the script stops, letting you fix the input file with the incompatibility. Afterwards, you can rerun the script to continue with the transition process.

### Before you begin

When running the script, you must specify the location of the CDH 5 Solr binaries using the `CDH_SOLR_HOME` environment variable. Solr binaries are located at `/opt/cloudera/parcels/CDH/lib/solr`.

For example:

```
export CDH_SOLR_HOME=/opt/cloudera/parcels/CDH/lib/solr
```

For information on `solr-upgrade.sh` command syntax and usage options, run the following command:

```
./solr-upgrade.sh help
```



**Warning:** Although the configuration transformation script addresses known incompatibilities, there might be incompatible configurations that are not detected by the script. Do not upgrade your production Cloudera Search services to Cloudera Runtime 7.1.1 or higher until you have tested the migrated configuration with Cloudera Runtime 7.1.1 or higher on a testing or development cluster.



## About this task

In this procedure you run `solr-upgrade.sh` with the `upgrade-metadata` option. It transforms configuration metadata to make it compatible with the target Solr version. The input of the script (`[***/SOLR/METADATA/INPUT/DIRECTORY***)` is the Solr configuration that you have downloaded from the ZooKeeper service of the source version.

The Solr configuration transition script, `solr-upgrade.sh`, is included with Cloudera Manager 7.1 agent software. This enables you to run the script after upgrading to Cloudera Manager 7.1, but before upgrading to Cloudera Runtime 7.1.1 or higher.

The script is located at `/opt/cloudera/cm/solr-upgrade/solr-upgrade.sh`.

## Procedure

1. Make sure that you run the script on a host that is assigned a Solr Server or Solr service Gateway role. Confirm that the `SOLR_ZK_ENSEMBLE` environment variable is set in `/etc/solr/conf/solr-env.sh`:

```
cat /etc/solr/conf/solr-env.sh
```

```
export SOLR_ZK_ENSEMBLE=[***/ZK01.EXAMPLE.COM:2181,ZK02.EXAMPLE.COM:2181/SOLR***/] \
export SENTRY_CONF_DIR=/etc/solr/[***/CONF.EXAMPLE.SOLR-1888]/sentry-conf
```

Replace `[***/ZK01.EXAMPLE.COM:2181,ZK02.EXAMPLE.COM:2181/SOLR***/]` and `[***/CONF.EXAMPLE.SOLR-1888]` with actual values valid in your environment.

2. Run the migration script:

```
/opt/cloudera/cm/solr-upgrade/solr-upgrade.sh upgrade-metadata \
-c [***/SOLR/METADATA/INPUT/DIRECTORY***/] \
-d [***/SOLR/METADATA/OUTPUT/DIRECTORY***/]
```

Replace `[***/SOLR/METADATA/INPUT/DIRECTORY***/]` and `[***/SOLR/METADATA/OUTPUT/DIRECTORY***/]` with actual values valid in your environment.

For example:

```
/opt/cloudera/cm/solr-upgrade/solr-upgrade.sh upgrade-metadata \
-c $HOME/cr7-solr-migration -d $HOME/cr7-migrated-solr-config
```

If you have enabled Kerberos, specify your JAAS configuration file by adding `--jaas [***/PATH/TO/SOLR/JAAS.CONF***/]` to the command.

For example:

```
/opt/cloudera/cm/solr-upgrade/solr-upgrade.sh \
--jaas [***/PATH/TO/SOLR/JAAS.CONF***/] \
upgrade-metadata -c $HOME/cr7-solr-migration -d $HOME/cr7-migrated-solr-config
```

The output directory does not only contain the migrated configuration files but several `*_validation.html` files for all the steps of the migration and a new `index_validation.html` file linking to these files.

The output contains the following logs at the beginning and end of the steps:

```
----- upgrading ...
----- upgrade successful for ...
```

If a step reports an error, the script stops at that point and the `METADATA UPGRADE SUCCESSFUL` message is missing. Check the output for the problematic step. If the script stops on error, the last step printed only has the upgrading but not the upgrade successful line.

- Fix the erroneous file in the input directory (`$HOME/cr7-solr-migration` in this example) and re-run the script.

Each time you run the script, the files in the output directory (`$HOME/cr7-migrated-solr-config` in this example) are overwritten. Repeat until the script outputs no incompatibilities.

If the script runs successfully, the last line of the output is the following:

```
METADATA UPGRADE SUCCESSFUL FOR METADATA IN [ ***/SOLR/METADATA/INPUT/
DIRECTORY*** ]
OUTPUT STORED IN [ ***SOLR/METADATA/OUTPUT/DIRECTORY*** ]
```

For example:

```
METADATA UPGRADE SUCCESSFUL FOR METADATA IN $HOME/cr7-solr-migration
OUTPUT STORED IN $HOME/cr7-migrated-solr-config
```

### Validate transitioned Solr configuration

The `solr-upgrade.sh` script includes a `validate-metadata` command that you can run against the migrated Solr configuration and metadata to make sure that they can be used to reinitialize the Solr service after the CDH cluster upgrade.

### About this task

The script performs a series of checks to make sure that:

- Required configuration files (such as `solr.xml`, `clusterstate.json`, and collection configuration sets) are present.
- The configuration files are compatible with the Solr version being upgraded to (Solr 8, in this case).

### Before you begin

#### Procedure

Validate the configuration by running the following script:

```
/opt/cloudera/cm/solr-upgrade/solr-upgrade.sh \
validate-metadata -c [ ***SOLR/METADATA/OUTPUT/DIRECTORY*** ]
```

Replace `[***SOLR/METADATA/OUTPUT/DIRECTORY***]` with the path to the directory containing the migrated Solr configuration.

If you have enabled Kerberos, specify your JAAS configuration file by adding `--jaas [***PATH/TO/SOR/JAAS.CONF***]` to the command.

For example:

```
/opt/cloudera/cm/solr-upgrade/solr-upgrade.sh validate-metadata \
-c $HOME/cr7-migrated-solr-config
```

### Results

If the validation is successful, the script outputs a message similar to the following:

```
Validation successful for metadata in /home/solruser/[ ***SOLR/METADATA/
OUTPUT/DIRECTORY*** ]
```

For example,

```
Validation successful for metadata in
/home/solruser/cr7-migrated-solr-config
```

If the validation fails, you can revisit the steps in *Transition the configuration*.

### Test transitioned Solr configuration on a Cloudera Runtime cluster

Copy the upgraded Solr configuration to a test or development CDP cluster to test it for incompatibilities not detected by the upgrade script before you initiate the upgrade on your production environment.

#### About this task



**Warning:** Although the configuration transformation script addresses known incompatibilities, there might be incompatible configurations that are not detected by the script. Do not upgrade your production Cloudera Search services to Cloudera Runtime 7.1.1 or higher until you have tested the migrated configuration with Cloudera Runtime 7.1.1 or higher on a testing or development cluster.

To test the migrated configuration on a Cloudera Runtime 7.1.1 or higher cluster, you can either provision a new host in your upgraded Cloudera Manager 7 environment and add a Cloudera Runtime 7.1.1 or higher cluster using that host, or you can upgrade a development or test cluster to Cloudera Runtime 7.1.1 or higher. You can then use the Cloudera Runtime 7.1.1 or higher cluster to test the migrated configuration to make sure that it works on Cloudera Runtime 7.1.1 or higher before upgrading your production environment.

#### Procedure

1. Create a TAR file of the configuration files in the `[***SOLR/METADATA/OUTPUT/DIRECTORY***]` (in this example `cr7-migrated-solr-config`) directory and copy them over to a Solr host on the test CDP cluster.

```
tar -czvf cr7-migrated-solr-config.tar.gz \
/home/solruser/cr7-migrated-solr-config
```

2. On the Solr host in the test CDP cluster, extract the TAR file created in the previous step.

```
tar -xzvf cr7-migrated-solr-config.tar.gz
```

3. Kinit as a solr user in the test CDP cluster:

```
kinit solr@EXAMPLE.COM
```

Replace `solr@EXAMPLE.COM` with your Kerberos realm name.

4. Run the following commands to create a Solr collection:

```
solrctl instancedir --create [***PATH/TO/MIGRATED/CONFIG**]/cr7-migrate
d-solr-config/configs/testcollection_config
```

```
solrctl collection --create testcollection_config
```

- Replace `[***CONFIG***]` with the name of the configuration (NOT the name of the collection) that you want to test.
- Replace `[***PATH/TO/MIGRATED/CONFIG***]` with the path to where you extracted `cr7-migrated-solr-config.tar.gz`.

If Kerberos is enabled and configured ZooKeeper Access Control Lists (ACLs), specify your JAAS configuration file by adding the `--jaas` parameter to the command. For example:

```
solrctl --jaas $HOME/SOLR-JAAS.CONF \
instancedir --create testcollection_config \
/[***PATH**]/cr7-migrated-solr-config/configs/testcollection_config
```

```
solrctl collection --create testcollection_config
```

5. Check the Solr web UI on the host where the above collections were created, and make sure the collection is created.
6. Repeat the above steps for all collections you want to transition to CDP.

## Preparing Hive for upgrade

Hive needs house cleaning before an upgrade. Some house cleaning tasks are mandatory to prevent upgrade failure. A smooth upgrade experience justifies spending the time and effort to clean up problems. Some problems prevent conversion of old Hive tables to Hive 3 ACID tables.

To prepare Hive for upgrade, you can perform the tasks below.

### Prerequisites

- Your Hive tables reside in HDP 2.6.5.x. Upgrades from HDP 3.x are not supported in this release of Cloudera Private Cloud Base.
- JDK is installed on the node running Hive Metastore.
- Hive Metastore is running and connected to the node where you will run the Hive pre-upgrade tool.
- If Hive metastore contains ACID tables, you have enabled ACID operations using Ambari Web or set the following Hive configuration properties to enable ACID. Failure to set these properties before upgrading Hive will result in corrupt or unreadable ACID tables.
  - `hive.txn.manager=org.apache.hadoop.hive.ql.lockmgr.DbTxnManager`
  - `hive.support.concurrency=true`
- Optionally, you shut down HiveServer2. Shutting down HiveServer2 is recommended, but not required, to prevent operations on ACID tables while the tool executes.
- You ensured the cluster has sufficient capacity to execute any compaction jobs that the upgrade tool might submit by setting `hive.compactor.worker.threads` to accommodate your data.
- You obtained permissions below to perform the steps to prepare Hive for upgrade.
  - Hive service user permissions, or all permissions to access Hive that Ranger provides
  - If you use Kerberos, a valid Kerberos ticket for starting Hive as the Hive service user
 

The Hive service user is usually the default hive user. If you don't know who the Hive service user is, go to the Ambari Web UI, and click Cluster Admin > Service Accounts, and then look for Hive User.
  - Capability to login as the HDFS superuser.
  - If you use Kerberos, you have become the HDFS superuser with a valid ticket.
- If you use Oracle as the backend database for Hive 1.x - Hive 3.x and the ojdbc7 JAR, replace this JAR with ojdbc6 JAR as described in the Cloudera Community ["Unable to start Hive Metastore during HDP upgrade"](#) article.

### Take a Mandatory Snapshot of Hive Tables

Taking a snapshot of Hive tables is mandatory before upgrading. You also need to keep track of how many tables you have before upgrading for comparison after upgrading.

### Procedure

1. In Ambari, go to Services/Hive/Configs, and check the value of `hive.metastore.warehouse.dir` to determine the location of the Hive warehouse, `/apps/hive/warehouse` by default.
2. On any node in the cluster, as the HDFS superuser, enable snapshots.

```
$ sudo su - hdfs
$ hdfs dfsadmin -allowSnapshot /apps/hive/warehouse
```

```
Allowing snapshot on /apps/hive/warehouse succeeded
```

3. Create a snapshot of the Hive warehouse.

```
$ hdfs dfs -createSnapshot /apps/hive/warehouse
```

Output includes the name and location of the snapshot.

```
Created snapshot /apps/hive/warehouse/.snapshot/s20181204-164645.898
```

4. Start Hive as a user who has SELECT privileges on the tables.

```
$ beeline beeline> !connect jdbc:hive2://
Enter username for jdbc:hive2://: hive
Enter password for jdbc:hive2://: *****
```

```
Connected to: Apache Hive (version 1.2.1000.2.6.5.0-292)
Driver: Hive JDBC (version 1.2.1000.2.6.5.0-292)
```

5. Identify all tables outside /apps/hive/warehouse/.

```
hive> USE my_database;
hive> SHOW TABLES;
```

6. Determine the location of each table using the DESCRIBE command. For example:

```
hive> DESCRIBE FORMATTED my_table partition (dt='20201130');
```

7. Create a snapshot of the directory shown in the location section of the output.

8. Repeat steps 5-7 for each database and its tables outside /apps/hive/warehouse/.

### Make Tables SparkSQL Compatible

Non-Acid, managed tables in ORC or in a Hive Native (but non-ORC) format that are owned by the POSIX user hive will not be SparkSQL-compatible after the upgrade unless you perform manual conversions.

#### About this task

If your table is a managed, non-ACID table, you can convert it to an external table using this procedure (recommended). After the upgrade, you can easily convert the external table to an ACID table, and then use the Hive Warehouse Connector to access the ACID table from Spark.

Take one of the following actions.

- Convert the tables to external Hive tables before the upgrade.
 

```
ALTER TABLE ... SET TBLPROPERTIES('EXTERNAL'='TRUE','external.table.purge'='true')
```
- Change the POSIX ownership to an owner other than hive.

You will need to convert managed, ACID v1 tables to external tables after the upgrade.

### Download the Pre-Upgrade Tool JAR for Compaction

You must download the software that you run later to detect tables for compaction and to perform that compaction.

#### Procedure

1. SSH into the host running the Hive Metastore.
2. Change to the /tmp directory.
3. Download the pre-upgrade tool JAR.

```
$ wget https://repo.hortonworks.com/repository/releases/org/apache/hive/hive-pre-upgrade/3.1.0.3.1.0.258-3/hive-pre-upgrade-3.1.0.3.1.0.258-3.jar
```

## Get a Kerberos Ticket

You must perform the recommended steps if you use Kerberos; otherwise, you can skip these steps.

### Procedure

1. Become the Hive service user.

For example, on Linux:

```
$ kinit -kt /etc/security/keytabs/hive.service.keytab hive/`hostname -f`
```

2. If you see the following error, modify the pre-upgrade tool command to include `-Djavax.security.auth.useSubjectCredsOnly=false`. The error is: `org.ietf.jgss.GSSEException: No valid credentials provided (Mechanism level: Failed to find any Kerberos tgt.`

For example, in step 4 of the procedure in [Preparing Hive for upgrade](#), modify the pre-upgrade tool command as follows:

```
$ $JAVA_HOME/bin/java -cp /usr/hdp/$STACK_VERSION/hive/lib/derby-10.10.2.0.jar:/usr/hdp/$STACK_VERSION/hive/lib/*:/usr/hdp/$STACK_VERSION/hadoop/*:/usr/hdp/$STACK_VERSION/hadoop/lib/*:/usr/hdp/$STACK_VERSION/hadoop-mapreduce/*:/usr/hdp/$STACK_VERSION/hadoop-mapreduce/lib/*:/usr/hdp/$STACK_VERSION/hadoop-hdfs/*:/usr/hdp/$STACK_VERSION/hadoop-hdfs/lib/*:/usr/hdp/$STACK_VERSION/hadoop/etc/hadoop/*:/tmp/hive-pre-upgrade-3.1.0.3.1.0.258-3.jar:/usr/hdp/$STACK_VERSION/hive/conf/server:/etc/hadoop/conf:/etc/hive/conf/ org.apache.hadoop.hive.upgrade.acid.PreUpgradeTool > {hive_log_dir}/pre_upgrade_{target_version}.log
```

For more information, see [Oracle Basic client server](#).

3. Perform the procedure for compacting Hive tables that follows.

## Run Compaction on Hive Tables

You must use the Hive pre-upgrade tool JAR you downloaded to run compaction on Hive tables. The pre-upgrade tool looks for files in an ACID table that contains update or delete events, and generates scripts to compact these tables. You can run the pre-upgrade tool command on the command line before or after upgrading Ambari. You do not actually use Ambari to run this command.

### About this task

Some transactional tables in HDP 2.6.5.x require a major compaction before upgrading the tables to Hive 3.0 in Cloudera Private Cloud Base. Running the Hive pre-upgrade tool identifies the tables that need such a compaction and provides scripts that you run to perform the compaction. Depending on the number of tables and partitions, and the amount of data involved, compactions might take a significant amount of time and resources. The script output of the pre-upgrade tool includes some heuristics that might help estimate the time required. If no script output is produced, no compaction is needed.

### Procedure

1. Log in as the hive user.

```
$ sudo su - hive
```

2. Export the JAVA\_HOME environment variable if necessary.

```
$ export JAVA_HOME=[ path to your installed JDK ]
```

3. Set STACK\_VERSION to the HDP version you are running.

```
$ export STACK_VERSION=`hdp-select status hive-server2 | awk '{ print $3; }'`
```

#### 4. Run the pre-upgrade tool command.

```
$ $JAVA_HOME/bin/java -cp /usr/hdp/$STACK_VERSION/hive/lib/derby-10.10.2
.0.jar:/usr/hdp/$STACK_VERSION/hive/lib/*:/usr/hdp/$STACK_VERSION/hadoop
/*:/usr/hdp/$STACK_VERSION/hadoop/lib/*:/usr/hdp/$STACK_VERSION/hadoop-m
apreduce/*:/usr/hdp/$STACK_VERSION/hadoop-mapreduce/lib/*:/usr/hdp/$STAC
K_VERSION/hadoop-hdfs/*:/usr/hdp/$STACK_VERSION/hadoop-hdfs/lib/*:/usr/h
dp/$STACK_VERSION/hadoop/etc/hadoop/*:/tmp/hive-pre-upgrade-3.1.0.3.1.0.
258-3.jar:/usr/hdp/$STACK_VERSION/hive/conf/conf.server:/etc/hadoop/conf
/;/etc/hive/conf/ org.apache.hadoop.hive.upgrade.acid.PreUpgradeTool > {
hive_log_dir}/pre_upgrade_{target_version}.log
```

The output indicates whether you need to perform compaction or not. In the /tmp directory, scripts named compacts\_nnnnnnnnnnnnnn.sql appear that contain ALTER statements for compacting tables.

```
ALTER TABLE default.t COMPACT 'major';
- Generated total of 1 compaction commands
- The total volume of data to be compacted is 0.001155MB
```

From the volume of data to be compacted, you can gauge how long the actual upgrade might take. If no scripts appear, a message in the output says you do not need to compact tables: ... org.apache.hadoop.hive.upgrade.acid.PreUpgradeTool - No compaction is necessary

#### 5. Check the logs on the Hive Metastore host for any errors.

- {hive\_log\_dir}/pre\_upgrade\_{target\_version}.log
- /tmp/hive/hive.log

#### 6. If there are no errors, proceed; otherwise, resolve the errors, and repeat this procedure.

#### 7. On the node where the Hive Metastore resides, log in as a user who has privileges to alter the Hive database.

#### 8. Start Beeline as the Hive service user.

```
$ beeline -u 'jdbc:hive2://<Metastore host name>:10000' -n hive
```

#### 9. On the Hive command line run the compaction script.

```
hive> !run /tmp/compacts_nnnnnnnnnnnnnn.sql
```

Output confirms that compaction is queued:

```
INFO : Compaction enqueued with id 3
...
```

### Save Hive Metastore by Dumping

You must dump the Hive metastore metadata immediately before upgrading. Dumping the metastore data to a file is critical to prevent data loss.

#### Procedure

##### 1. On the node where the database you use for Hive Metastore resides, dump Hive Metastore metadata before upgrading to HDP.

For example, in MySQL, dump each database as follows:

```
mysqldump <hive_db_schema_name> > </path/to/dump_file>
```

In Postgres, dump each database as follows:

```
pg_dump --opt <metastore_db_name> > /path/to/dump_file
```

- Proceed to upgrade HDP, assuming no Hive update, delete, or merge occurred after compaction; otherwise, repeat the compaction and Hive Metastore dump procedures, and then upgrade HDP.

### Capture Information about Multiple HiveServers

If you have multiple HiveServer (HS2) pods that either fence business units from each other or offer doas impersonation types to aid a transition, you must configure Cloudera Manager accordingly after the upgrade. You need to note the configuration information from Ambari to use for this configuration task after the upgrade.

#### Procedure

- Make a note of the Ambari Config Group settings for each HS2 group, including the edge nodes and connected endpoints.
- Also, make a note of the Config Groups you use for other purposes, if there are any.
  - Run as end user instead of Hive user (hive.server2.enable.doAs)
  - ZooKeeper Namespace (hive.server2.zookeeper.namespace)
  - Transport Mode (hive.server2.transport.mode)

```
Config
Run as end user instead of Hive user (hive.server2.enable.doAs)

ZooKeeper Namespace (hive.server2.zookeeper.namespace)

Transport Mode (hive.server2.transport.mode)
Value _____
Cloudera Manager _____
```

### Hive Pre-Upgrade Tool Command Help

The pre-upgrade tool command supports a few options that you can consider using while running it. Take a look at the output of the help command to see all the options.

#### -d <database regex>

Regular expressions to define databases. Limits the scope of the tool.

#### -execute

Do not use this option unless you want to run the pre-upgrade tool command in Ambari instead of on the Beeline command line. Using Beeline instead of this option is strongly recommended. This option automatically executes the equivalent of the generated commands.

#### -location <path>

Use this option to specify the location to write the scripts generated by the pre-upgrade tool.

#### -t <table regex>

Regular expressions to define tables. Limits the scope of the tool.

#### -tn

Number of the thread pool.

If there is a large number of tables that need to be compacted, run the tool in a parallel mode with the -tn option. This controls the number of threads used to scan tables.



**--help**

You can append `--help` to the command to see all command options. For example:

```
java -Djavax.security.auth.useSubjectCredsOnly=false -cp /usr/hdp/$STACK_VERSION/hive/lib/derby-10.10.2.0.jar:/usr/hdp/$STACK_VERSION/hive/lib/*:/usr/hdp/$STACK_VERSION/hadoop/*:/usr/hdp/$STACK_VERSION/hadoop/lib/*:/usr/hdp/$STACK_VERSION/hadoop-mapreduce/*:/usr/hdp/$STACK_VERSION/hadoop-mapreduce/lib/*:/usr/hdp/$STACK_VERSION/hadoop-hdfs/*:/usr/hdp/$STACK_VERSION/hadoop-hdfs/lib/*:/usr/hdp/$STACK_VERSION/hadoop/etc/hadoop/*:/tmp/hive-pre-upgrade-3.1.0.3.1.0.258-3.jar:/usr/hdp/$STACK_VERSION/hive/conf/conf.server:/etc/hadoop/conf:/etc/hive/conf/org.apache.hadoop.hive.upgrade.acid.PreUpgradeTool --help
```

**Preparing the backend HMS database for upgrade**

Learn how you can prevent upgrade failures when you have materialized views or MSSQL indexed views created on top of Hive backend schemas, such as `SYS` or `INFORMATION_SCHEMA` tables.

**About this task**

If you have created any materialized views or MSSQL indexed views on Hive `SYS` or `INFORMATION_SCHEMA` tables, your upgrade process can fail when the upgrade SQL statements are trying to drop these tables.

You must drop the materialized views before performing an upgrade and then recreate the views after the upgrade process is complete.



**Important:** Cloudera recommends that you do not create any materialized views or indexed views on the `SYS` and `INFORMATION_SCHEMA` tables. However, if you have already created the views, then follow the steps provided in this topic to identify such views and drop them before upgrading the cluster.

**Before you begin**

You must back up the Hive metastore (HMS) backend database before dropping the materialized views.

**Procedure**

1. Start a Hive Beeline session and run the following query to identify materialized views that are created on top of the `SYS` and `INFORMATION_SCHEMA` tables:

```
SELECT DISTINCT d.DB_LOCATION_URI, d.NAME, t.TBL_NAME, t.TBL_TYPE, t.OWNER, t.VIEW_EXPANDED_TEXT
FROM sys.TBLS t
    INNER JOIN sys.DBS d ON t.DB_ID = d.DB_ID
    INNER JOIN sys.MV_CREATION_METADATA mv ON mv.TBL_NAME = t.TBL_NAME
    INNER JOIN sys.MV_TABLES_USED tu ON mv.MV_CREATION_METADATA_ID = tu.MV_CREATION_METADATA_ID
WHERE tu.TBL_ID IN (SELECT distinct t.TBL_ID
    FROM sys.MV_CREATION_METADATA mv
        INNER JOIN sys.MV_TABLES_USED tu ON mv.MV_CREATION_METADATA_ID = tu.MV_CREATION_METADATA_ID
        INNER JOIN sys.TBLS t ON tu.TBL_ID = t.TBL_ID
        INNER JOIN sys.DBS d ON t.DB_ID = d.DB_ID
        WHERE lower(d.NAME) IN ('sys', 'information_schema'))
AND upper(t.TBL_TYPE) = 'MATERIALIZED_VIEW';
```

2. If the query returns any materialized views, drop each view using the `DROP` statement.
 

```
DROP MATERIALIZED VIEW [db_name.]materialized_view_name;
```
3. Upgrade the cluster and recreate the views after the upgrade process is complete.

## Preparing Kafka for upgrade

To successfully upgrade Kafka, you must set the protocol version to match the protocol version used by the brokers and clients.

1. In Ambari Configs, update the following properties in the `server.properties` file:

- `inter.broker.protocol.version = current_Kafka_version`
- `log.message.format.version = current_Kafka_version`



**Note:** `current_Kafka_version` is the version of Kafka you are currently running. For example, 1.0.

Review the *HDP Release Notes* for the Apache Kafka version in use with your HDP installation.

2. Save your changes.

### Extract Kafka broker ID

You must extract the Kafka broker IDs before you upgrade the HDP cluster from 2.6.5 to HDP intermediate bits.

You must extract the Kafka broker IDs manually from the HDP 2.6.5 cluster. When you migrate to Cloudera Manager, you must manually enter the broker IDs to the real values in Kafka.

#### Procedure 1

To extract the broker IDs from the HDP 2.6.5 cluster, use the following commands: This procedure helps you only if the Kafka service is not stopped.

- For Unsecure environments: `sh /usr/hdp/current/kafka-broker/bin/kafka-broker-api-versions.sh --bootstrap-server <host:port> | grep "id" | awk -F'[=:]' '{print $1,$3}' | awk '{print $1,$2}' > /tmp/kafka-broker-ids.ini`
- For Secure environments: `sh /usr/hdp/current/kafka-broker/bin/kafka-broker-api-versions.sh --bootstrap-server <host:port> --command-config <client.config> | grep "id" | awk -F'[=:]' '{print $1,$3}' | awk '{print $1,$2}' > /tmp/kafka-broker-ids.ini`
- For example, `sh /usr/hdp/current/kafka-broker/bin/kafka-broker-api-versions.sh --bootstrap-server c6111-node2.coelab.cloudera.com:6667 --command-config <client.config> | grep "id" | awk -F'[=:]' '{print $1,$3}' | awk '{print $1,$2}' > /tmp/kafka-broker-ids.ini`

where `<client.config>` contains environment specific Kafka client security configuration for connecting the source Kafka cluster. For more information on configuring Kafka client, see [Kafka documentation](#) and [Kafka client configuration](#).



**Note:** This command generates the `kafka-broker-ids.ini` file. The `kafka-broker-ids.ini` file must be manually copied to `$am2cm-1.1.1.0-xx/conf/` path (if you are upgrading to CDP Private Cloud Base 7.1.6) or `$am2cm-1.1.0.0-xx/conf/` path (if you are upgrading to CDP Private Cloud Base 7.1.4) before running the AM2CM tool.

#### Procedure 2

1. Create `kafka-broker-ids.ini` file
2. Navigate to each Kafka broker host `> $log.dirs/meta.properties`
3. Pick up the `broker.id` value.
4. Copy the hostname `broker.id` to `kafka-broker-ids.ini` file

An example of the file format:

```
ctr-e153-xxxx-xxxx71.cloudera.site 1001
ctr-e153-xxxx-xxxx72.cloudera.site 1002
ctr-e153-xxxx-xxxx73.cloudera.site 1003
```

## Preparing Zeppelin for upgrade

Before upgrading HDP to CDP, you must copy the Zeppelin notebooks and configuration files from HDFS to the Zeppelin server host filesystem.



**Important: Deprecation notice for Zeppelin:** Zeppelin is deprecated in Cloudera Runtime 7.1.9 and 7.2.18. For more information, see the deprecation notices in the corresponding Cloudera Runtime release notes.

### Procedure

1. On the Zeppelin server host, create the `/var/lib/zeppelin/conf` and `/var/lib/zeppelin/notebook` directories:

```
sudo -u zeppelin mkdir -p /var/lib/zeppelin/conf /var/lib/zeppelin/notebook
```

2. If you are using Kerberos, authenticate as the `zeppelin` user first. Otherwise, run the commands as the `zeppelin` user by prefixing the commands with `sudo -u zeppelin`:

```
kinit zeppelin
```

3. Copy all notebooks from HDFS to the `/var/lib/zeppelin/notebook` directory. If you are using Kerberos, authenticate as the `zeppelin` user first:

```
kinit zeppelin
hdfs dfs -get "/user/zeppelin/notebook/*" /var/lib/zeppelin/notebook
```

4. Copy the `notebook-authorization.json` files from HDFS to the `/var/lib/zeppelin/conf` directory:

```
hdfs dfs -get /user/zeppelin/conf/notebook-authorization.json /var/lib/zeppelin/conf
```

### What to do next

Continue with the upgrade.

## Perform the HDP upgrade

After you have performed the required pre-upgrade steps and backed up all the relevant data, you can proceed with the process to upgrade your cluster to HDP intermediate bits.

### Procedure

1. Log in to Ambari.
2. Browse to `Cluster Admin Stack and Versions`.
3. Click the `Versions` tab.

The registered and installed target HDP version displays an `Upgrade` button.

4. Click `Upgrade` on the target version.

Based on your current HDP version and the target HDP version, Ambari performs a set of prerequisite checks to determine if you can perform a rolling or an express upgrade. A dialog displays the options available.

5. Select the Express Upgrade method. (Only supported method for upgrading from HDP 2.6.5.x to HDP intermediate bits.) Advanced options are also available.
  - Skip all Service Check failures:
 

Ambari automatically skips any Service Check failures and completes the task without requiring user actions. After all the Service Checks have run in a task, you see a summary of the failures and options to continue the upgrade or pause.
  - Skip all Slave Component failures
 

Ambari automatically skips any Secondary Component failures and completes the task of upgrading Secondary components without requiring user actions. After all the Secondary Components have been upgraded, you see a summary of the failures and options to continue the upgrade or pause.
6. Click Proceed.

## Perform express upgrade

You must use the Express Upgrade method to upgrade your cluster from HDP 2.6.5.x to HDP intermediate bits.

### Procedure

1. Ambari checks that your cluster meets prerequisites. A dialog displays the results:
  - If any required prerequisites are not met, the result displays an error.
 

You cannot proceed with the upgrade until you make the appropriate corrections and return to Perform Upgrade again.
  - If any optional prerequisites are not met, the result displays a warning.
 

You may proceed with the upgrade.
  - Ambari displays a list of configuration changes that occur during the upgrade.
2. When the prerequisite checks complete, the upgrade starts. The time required to perform the upgrade depends on many factors. As part of the upgrade process, each component in the cluster restarts in a serial fashion. The stop/start times contribute to the total upgrade time.
3. The upgrade process includes the following stages. Some stages require that you complete an action during normal operation. If any stage fails, the upgrade stops and prompts you for action.

| Stage                                   | Description                                                                                                                                                                                                                               | Action required                                       |
|-----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------|
| Prepare Upgrade                         | You should stop all apps on YARN queues, and deactivate & kill all running Storm topologies. If you have any applications running on Slider, it is recommended to stop them, since Slider will be removed as part of the Upgrade process. | Perform the actions to prepare for the upgrade.       |
| Stop Components for High-Level Services | This will stop all components for High-Level Services. This includes all master components except those of HDFS, HBase, ZooKeeper and Ranger.                                                                                             | None                                                  |
| Perform Backups                         | This step prompts you to confirm that you have taken proper backups before proceeding.                                                                                                                                                    | You must acknowledge the prompt for database backups. |
| Stop Components for Core Service        | Stops all components with HDFS, HBase, ZooKeeper and Ranger.                                                                                                                                                                              | None                                                  |
| Update Target Repository                | Updates the stack version in Ambari to the target version.                                                                                                                                                                                | None                                                  |
| Add New YARN Components                 | In Cloudera Private Cloud Base, there are two new YARN components: YARN Registry DNS, and Timeline Reader. Both of these components are automatically added to the cluster.                                                               | None                                                  |

|                            |                                                                                                                                                                     |                                                                                                                                                                                                          |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Update Service Configs     | Updates (i.e. transfers or replaces) any configurations that are necessary for the upgrade.                                                                         | None                                                                                                                                                                                                     |
| Restart Components         | Restarts all core components such as ZooKeeper, Ranger, HDFS, YARN, MapReduce2 and various Clients (Tez, Pig, Sqoop).                                               | In an SSO-enabled cluster: If the cookie is lost/session is expired:client should use local login to access the Ambari and proceed further.<br>For example: <ambari_host:ambari_port>/#/login/local      |
| All Service Checks         | All Service Checks are performed against the cluster.                                                                                                               | Any service check that fails prompts you to Ignore and Continue, Downgrade or Retry. If you selected the Skip all Service Check failures option, you are only prompted when all Service Checks complete. |
| Restart Components         | Restarts the remaining components such as Oozie, Hive, Spark2 and others.                                                                                           | None                                                                                                                                                                                                     |
| Set Version on All Hosts   | Sets the HDP version on all hosts to the target HDP version.                                                                                                        | None                                                                                                                                                                                                     |
| Finalize Upgrade Pre-Check | Checks if any hosts were not upgraded, either because the host was in Maintenance Mode, or one or more components on the host failed to upgrade (and were skipped). | Click the list that displays # hosts for details on the hosts (and their components) that are not upgraded. You can Pause Upgrade, delete the hosts and return to finalize.                              |
| Finalize Upgrade           | The component upgrades are complete. You are presented the option to Finalize, which when selected, completes the upgrade process + saves the cluster state.        | When prompted to remove HBase snapshots, look for snapshots matching \$TABLE-ru-\$TIMESTAMP<br>Prompted to Finalize or Finalized Later or Downgrade.                                                     |

- When the upgrade stages complete, you can choose to Finalize the upgrade, or to Finalize Later. Finalizing later gives you a chance to perform more validation on the cluster.



**Note:** If you choose to finalize later, both versions are listed on the Stack and Versions tab with the starting version displaying as Current. It is not until you finalize that Ambari makes the target version the current version. Also, until you finalize, you will not be able to perform operational changes to the cluster (such as move components, change configurations, and so on).

- Click Finalize to complete the express upgrade process.

## Post-HDP-upgrade tasks

After the HDP cluster is upgraded to HDP intermediate bits, ensure that you check the cluster settings, migration, and restore settings are correct.

## Update Ranger passwords

Ranger password validation has been updated in HDP intermediate bits

conform to these new password policies, the following Ranger passwords need to be updated to ensure that they have at least 8 characters with minimum one alphabet and one numeric. These passwords cannot contain the following special characters: " ' \ `

- Ranger Admin

The following new passwords need to be populated with valid passwords that also conform to the password policy:

- Ranger Usersync User's Password
- Ranger Tagsync User's Password
- Ranger KMS Keyadmin User's Password

## Atlas Migration and HBase Hook settings

There are specific settings for the Atlas migration that you must update after the cluster upgrade is completed.

### About this task

After the upgrade to HDP intermediate bits, you must ensure that the Atlas settings are unchanged. The `atlas.migration.data.filename` property will continue to point to a non-existent path. For example, `atlas.migration.data.filename=/tmp/non-existent-path`.

### Procedure



**Note:** You must add Atlas Gateway role to all hosts which have the hook-based services like Hive / Hive on Tez / Impala / HBase / Spark / Kafka.

You must configure the settings for the Atlas migration so that metadata can be imported into Atlas. You must remove the configuration after the migration completes.

- After the migration import is complete, return to the `atlas.migration.data.filename` property and click Remove to remove it. Restart Atlas when prompted by Ambari.

When upgrading to HDP intermediate bits from earlier versions, you might need to manually enable the HBase hook.

- On the Ambari dashboard, select `HBaseConfigsAdvanced hbase-env`, then select `Enable Atlas Hook`. Click Save, then restart HBase and any other services that require a restart.

### Related Information

[Configure Atlas file-based authentication](#)

## Ambari Metrics and LogSearch

Take the Ambari Metrics and Log Search Services out of Maintenance Mode by choosing `Actions > Turn Off Maintenance Mode` from each Service page.

## Ambari infra-migrate and restore

Follow the steps below to restore the data previously backed up:

### Procedure

1. SSH to the host where the `migrationConfigGenerator.py` was run prior to the HDP Upgrade. This will be from one of your Ambari Infra Solr instances. Ensure you are in the current working directory containing the `ambari_solr_migration.ini` file.
2. Export the variable used to hold the path to the ini file.

```
export CONFIG_INI_LOCATION=ambari_solr_migration.ini
```

3. Migrate the data to ensure it's in the right format to import into Solr 7. Note that this script can take a long time to run depending on the volume of data backed up. It is recommended to run this script using the `nohup` command.

```
nohup /usr/lib/ambari-infra-solr-client/ambariSolrMigration.sh --ini-file  
\ $CONFIG_INI_LOCATION --mode migrate-restore
```

4. Re-index the migrated data into your current collections so the backed up data is visible in all of the tools using the Infra Solr instances. Note that this script can take a long time to run depending on the volume of data backed up. It is recommended to run this script using the `nohup` command.

```
nohup /usr/lib/ambari-infra-solr-client/ambariSolrMigration.sh --ini-file  
\ $CONFIG_INI_LOCATION --mode transport
```

## Upload HDFS entity information

You must also configure the property below to update the values of the Ranger tag store only if you are using Ranger Tagsync.

Property: `ranger.tagsync.atlas.hdfs.instance.c11.ranger.service`



**Note:** `c11` refers to cluster name.

Preconditions:

- If you have Tagsync configured previously, then manually set the value of the property to `cm_hdfs`.
- Not already set. Default service name is `c11_hadoop`.
- This property is applicable only if you are upgrading from HDP to CDP Private Cloud Base. If you are on a completely new cluster, you can ignore this property.

## Custom Spark SQL Warehouse Directory

If required, you must configure the Spark SQL warehouse location in the upgraded cluster.

If Applicable: Configure Custom Spark SQL Warehouse Directory

Since Spark 2.0.0, Spark references `spark.sql.warehouse.dir` as the default Spark SQL Warehouse location. In HDP 2.6.5.x, the standard value for the `spark.sql.warehouse.dir` property is `/apps/hive/warehouse`. On an Ambari cluster with HDP 2.6.5.x, this property must be set manually in both Custom `spark2-defaults` and Custom `spark2-thrift-sparkconf`. When upgrading from HDP 2.6.5.x to HDP intermediate bits, Ambari migrates the `spark.sql.warehouse.dir` property to Advanced `spark2-defaults` and Advanced `spark2-thrift-sparkconf`, and changes the value of this property to `/apps/spark/warehouse`. This is done to accommodate the separate Spark and Hive catalogs introduced in the HDP intermediate bits. If you used a custom setting for the `spark.sql.warehouse.dir` property in HDP 2.6.5.x, the Ambari upgrade to HDP intermediate bits ignores the custom setting and sets the value of the `spark.sql.warehouse.dir` property to `/apps/spark/warehouse` in both Advanced `spark2-defaults` and Advanced `spark2-thrift-sparkconf`.

If you want to use a custom Spark SQL warehouse after upgrading to the HDP intermediate bits, select `Spark2 > Configs`, then use `Add Property` in Advanced `spark2-defaults` and Advanced `spark2-thrift-sparkconf` to update the value of the `spark.sql.warehouse.dir` property with the custom setting.

If Hive is configured to run as a custom user, you must change the ownership of your new Hive warehouse directory to be owned by that custom user. `hdfs dfs -chown cstm-hive /warehouse/tablespace/managed/hive/`

## Hive post-HDP-upgrade tasks

You need to perform several critical tasks after upgrading to Cloudera Private Cloud Base, such as checking table locations and correcting locations of files. Before beginning these tasks, you need a little information about Hive changes to ACID properties.

A brief description of storage formats is provided as background information for understanding how the upgrade process relocates tables.

Changes include the location of tables, the HDFS path to the Hive warehouse, file ownership, table types, and ACID-compliance. You need to understand the following changes that occur during the upgrade before performing the post-upgrade tasks.

### Hive Changes to ACID Properties

Hive 2.x and 3.x have 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 is ACID v1. In Hive 3.x, the mature version of ACID is ACID v2, which is the default table type in HDP Cloudera Private Cloud Base.

### Native and Non-native Storage Formats in Hive

Storage formats are a factor in upgrade changes to table types. No change has occurred to storage formats in Cloudera Private Cloud Base, but storage formats determine table changes that occur during the upgrade. Hive 2.x and 3.x supports the following Hadoop 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`

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.

### Checking and correcting Hive table locations

As a Data Engineer, you need to understand the relocation of files after the upgrade process. The file type and other factors affect the relocation during the upgrade.

### About this task

The upgrade process changes table types in some cases. The following table compares Hive table types and ACID operations before and after upgrading. The ownership of the Hive table file is a factor in determining table types and ACID operations after the upgrade.

Table 4.1. Before and After Upgrading Table Type Comparison

| Before Upgrading |         |                      |                                 | After Upgrading             |         |
|------------------|---------|----------------------|---------------------------------|-----------------------------|---------|
| 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      |

\* See [Dropping an External Table Along with the Data](#).

\*\* Not SparkSQL-compatible

If you had external files before the upgrade, the upgrade process carries the external files over to CDP after upgrading with no change in location. The external files continue to reside in the `/apps/hive/warehouse` directory.

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 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 directory is in a different encryption zone than the new warehouse directory.

The `/apps/hive/warehouse` directory, which is the location of the Hive 2.x warehouse before upgrading, might or might not exist after upgrading.

### Procedure

1. Check the `/apps/hive/warehouse` directory for files that do not belong there after upgrading.

Files that do not belong in `/apps/hive/warehouse` are files described in the table above as managed files after upgrading. The upgrade process should have moved the managed files to the new `/warehouse/tablespace/managed/hive/warehouse` directory.

2. Check that the upgrade process moved managed files to `/warehouse/tablespace/managed/hive`.
3. Check that Hive places any new external tables you create after upgrading in `/warehouse/tablespace/external/hive`.

### Related Information

[Preventing SparkSQL incompatibility](#)

#### Preventing SparkSQL incompatibility

You need to be aware of two SparkSQL incompatibilities and how to work around these problems. The upgrade process converted all CDH Hive tables to external tables, however, if you moved managed, non-ACID tables preventing conversion to external tables, these are not compatible with native SparkSQL. Also, you might encounter a problem reading Hive 2 external ORC tables from Spark.

### Related Information

[Checking and correcting Hive table locations](#)

#### Managed, non-ACID table problem

#### About this task

Consider using either one of the following options:

- Convert ACID tables to external tables after the upgrade.
- Use the Hive Warehouse Connector.

Follow these steps to create a new external table using Hive 3 and migrate the data from the managed table to the new table.

### Procedure

1. Run the `SHOW CREATE TABLE` statement on the original table to get the full definition of the table.

```
SHOW CREATE TABLE <tablename>;
```

2. Rename the managed table to `*_old`.
3. Migrate data from `*_old` to `<new>` external table using the original name in the historical, or the default, location (`/warehouse/tablespace/external/hive/<?>.db/<tablename>`).

```
CREATE EXTERNAL TABLE new_t AS SELECT * FROM old_t;
```

#### Reading a Hive external table in ORC from Spark

#### About this task

Unlike the Hive ORC reader, which supports recursive directory reads, the Spark native ORC reader does not support recursive directory reads of Hive tables. Reading Hive 2.x tables is a problem under both of the following circumstances:

- You created the table using Hive CTAS (create table as select).

- One of more selected tables included UNION ALL.

When creating a table under these circumstances, subdirectories are named /1 /2 /3. Subdirectories do not include the HIVE\_UNION\_SUBDIR\_ prefix as Hive 3-created tables do. You cannot read these tables from Spark if the tables are in ORC format.

The following workaround configures Spark to overcome this problem.

### Procedure

1. If you already started the Spark shell, quit the shell.  
You cannot perform workaround configuration for the session if the session is already started.
2. Start the Spark shell with `convertMetastoreOrc` disabled.  
For example:

```
spark-shell ...
--conf spark.sql.hive.convertMetastoreOrc=false
```

### Correct Hive File Locations

Contacting Cloudera Support for help to correct file location problems is highly recommended. Alternatively, you can perform correct file locations of managed files yourself.

### Procedure

1. Log in as the HDFS superuser and if you use Kerberos, you need to become the superuser with a valid ticket.

```
$ sudo su - hdfs
```

2. Start Hive in Ambari 7.1.x.x.
3. On a node in your cluster, start Beeline in the background and a Hive shell in the foreground.

```
$ hive
```

4. Change the location of the database and the table from the old location to the new location.

```
hive> ALTER DATABASE tpcds_bin_partitioned_orc_10 SET LOCATION
'hdfs://ns1/apps/hive/warehouse/tpcds_bin_partitioned_orc_10.db';
hive> ALTER TABLE tpcds_bin_partitioned_orc_10.store_sales SET LOCATION
'hdfs://ns1/apps/hive/warehouse/tpcds_bin_partitioned_orc_10.db/store_sales';
```

5. On the Hive Metastore node, log in as the HDFS superuser.
6. Set `STACK_VERSION` to the HDP version you are running.

```
$ export STACK_VERSION=`hdp-select status hive-server2 | awk '{ print $3
; }'`
```

7. Run the script for strict Hive migration.

```
$ /usr/hdp/$STACK_VERSION/hive/bin/hive --config /etc/hive/conf
--service strictmanagedmigration
--hiveconf hive.strict.managed.tables=true -m automatic
--modifyManagedTables --oldWarehouseRoot /apps/hive/warehouse
```

### Handle Missing Table or Partition Locations

You need to identify missing table or partition locations, or both, to prevent upgrade failure. If the table and partition locations do not exist in the file system, you must either create a replacement partition directory (recommended) or drop the table and partition.

**About this task**

You perform this step if you did not modify the HSMM process to expedite the Hive upgrade.

**Procedure**

Ensure the table and partition locations exist on the file system. If these locations don't exist either create a replacement partition directory (recommended) or drop the table and partition.

**Managed Table Location Mapping**

A managed table location must map to one managed table only. If multiple managed tables point to the same location, upgrade problems occur.

**Check SERDE Definitions and Availability**

Ensure correct Serde definitions and a reference to a SERDE exists to ensure a successful upgrade.

**About this task**

You perform this step if you do not modify the HSMM process for expediting the Hive upgrade.

**Procedure**

1. Check Serde definitions for correctness and check for SERDE availability.
2. Correct any problems found as follows:
  - Remove the table having the problematic SERDE.
  - Ensure the SERDE is available during the upgrade, so the table can be evaluated.

**Verify Zeppelin settings in Ambari**

After upgrading Ambari and HDP, make sure that the Zeppelin settings reference the local filesystem paths instead of HDFS.



**Important: Deprecation notice for Zeppelin:** Zeppelin is deprecated in Cloudera Runtime 7.1.9 and 7.2.18. For more information, see the deprecation notices in the corresponding Cloudera Runtime release notes.

**Procedure**

1. In the Zeppelin configuration page, search for and verify the following settings:

```
zeppelin.notebook.storage = org.apache.zeppelin.notebook.repo.FileSystem
NotebookRepo
zeppelin.config.fs.dir = file:///var/lib/zeppelin/conf
zeppelin.notebook.dir = file:///var/lib/zeppelin/notebook
```

2. If the settings are missing or incorrect, add or update them to the above values.
3. In the Zeppelin UI, check the Interpreters page for any duplicate entries. If any duplicates exist:
  - a) Backup and then delete the interpreter.json file from HDFS (/user/zeppelin/conf/interpreter.json) and from the local filesystem on the Zeppelin server host (/var/lib/zeppelin/conf/interpreter.json).
  - b) Restart the Zeppelin service in Ambari.
  - c) In the Zeppelin UI, confirm that the duplicate entries no longer exist.
4. If any custom interpreter settings are missing, add them again from the Interpreters page.
5. Verify that your existing notebooks are still available.

**Related Information**

[Deprecation notice for Zeppelin](#)

## Search post-HDP-upgrade tasks

After upgrading your HDP cluster to Cloudera Runtime, you need to add a Solr service to your CDP cluster where you can recreate the collections you transitioned from HDP Search. Once done, you can proceed with reindexing data to the transitioned collections.

### Before you begin

- HDP is upgraded to Cloudera Runtime 7.1.1 or higher.
- The management platform is transitioned from Ambari to Cloudera Manager.

### Procedure

1. In the Cloudera Manager Admin Console, add Solr service to your CDP cluster.  
For more information, see [Adding a Service](#).
2. In the Cloudera Manager Admin Console, go to Solr service Configuration .
3. In the Search field, type upgrade to filter the configuration parameters.
4. For the Solr Server for Upgrade property, select one of your Solr Server hosts. When you run the migration script, make sure to run it on this host.
5. For the Upgrade Backup Directory property, specify a directory in HDFS. This directory must exist and be writable by the Solr service superuser (solr by default). Examples in these procedures use /cr7-solr-upgrade/backup for this HDFS directory.
6. For the Upgrade Metadata Directory property, specify a directory on the local filesystem of the host you selected as the Solr Server for Upgrade. When you run the migration script, make sure that you copy the migrated configuration to this directory. This directory must also be writable by the Solr service superuser. Examples in these procedures use /cr7-solr-metadata/migrated-config for this local directory.
7. Enter a Reason for change and then click Save Changes to commit the changes.
8. Copy the transitioned configuration metadata to the Solr server host designated as Solr Server for Upgrade.
9. Reinitialize Solr state for upgrade:
  - a) In the Cloudera Manager Admin Console, go to Solr service .
  - b) Select Actions Reinitialize Solr State for Upgrade and click Reinitialize Solr State for Upgrade to confirm.
10. Bootstrap the Solr configuration:
  - a) In the Cloudera Manager Admin Console, go to Solr service .
  - b) Select Actions Bootstrap Solr Configuration and click Bootstrap Solr Configuration to confirm.
11. Start the Solr service:
  - a) In the Cloudera Manager Admin Console, go to Solr service .
  - b) Select Actions Start .
12. Bootstrap the Solr collections:
  - a) In the Cloudera Manager Admin Console, go to Solr service .
  - b) Select Actions Bootstrap Solr Collections and click Bootstrap Solr Collections to confirm.

### What to do next

1. [Recreate aliases](#), if you were using them.
2. [Reindex Solr collections](#).

## Backup Infra Solr collections

Backup the Infra solr collections from the Ambari managed Infra Solr service and restore them on the Infra Solr service managed by Cloudera Manager. When you are creating a backup of the Infra solr data, ensure that only Infra Solr and Zookeeper services are started and other services are stopped. Similarly, when you are restoring the data on Infra Solr managed by Cloudera Manager, you must start HDFS and Zookeeper only. This ensures no operations are performed by services which can push data to Infra Solr.

### About this task

To create a backup of the Infra solr collections from Infra Solr service, perform the following steps:

### Procedure

1. The assumption is /opt/solrdata is the work directory which must pre-exist on both local file system and HDFS.



**Note:** If the Solr cloud is a multi node, then /opt/solrdata should be a common mount accessible from all the Solr server hosts.

```
mkdir -p /opt/solrdata/
kinit -kt hdfs.keytab hdfs; hdfs dfs -mkdir -p /opt/solrdata
```

2. Set the path for the ini configuration file.

```
CONFIG_INI_LOCATION=/opt/solrdata/ambari_solr_migration.ini
```

3. For the backup of data from Infra Solr to HDFS:

- a) Generate the ini configuration file which is used by the migration helper.

```
/usr/bin/python /usr/lib/ambari-infra-solr-client/migrationConfigGenerator.py --ini-file $CONFIG_INI_LOCATION --host <Ambari Server Host> --port <Ambari Server Port> --cluster <Cluster-Name> --username <Ambari admin username> --password <Ambari admin password> --backup-base-path=/opt/solrdata --hdfs-base-path=/opt/solrdata --java-home <JAVA JDK HOME which is also used by Ambari > -s < use -s if TLS is enabled for Ambari >
```

- b) Backup collections:

```
/usr/lib/ambari-infra-solr-client/migrationHelper.py --ini-file $CONFIG_INI_LOCATION --action backup
```

- c) Copy the backup to HDFS:

```
/usr/lib/ambari-infra-solr-client/migrationHelper.py --ini-file $CONFIG_INI_LOCATION --action copy-to-hdfs
```

You can now use the AM2CM tool to migrate from the Ambari-managed HDP cluster to the Cloudera Manager-managed CDP cluster.

## Troubleshooting the HDP upgrade

In the event of a problem, contacting Cloudera Support is highly recommended. Alternatively, you can perform these troubleshooting procedures.

### Hive Metastore corrupt

If the Hive Metastore is corrupt, you can perform specific steps to restore it.

### Procedure

1. On the node where the database for Hive Metastore resides, create databases you want to restore. For example:

```
$ mysql -u hiveuser -p -e "create database <hive_db_schema_name>;"
```

2. Restore each Metastore database from the dump you created. For example:

```
$ mysql -u hiveuser -p <hive_db_schema_name> < </path/to/dump_file>
```

3. Reconfigure Hive Metastore if necessary. Reconfiguration might be necessary if the upgrade fails. Contacting Cloudera Support for help with reconfiguration is recommended. Alternatively, in intermediate bits, set `key=value` commands on the command line to configure Hive Metastore.

## Missing Hive tables

In case you have lost the Hive tables, you can recover them by performing a few steps.

### Procedure

1. Login as Superuser HDFS. For example:

```
$ sudo su - hdfs
```

2. On HDFS, read the snapshots you saved of your table data. For example:

```
$ hdfs dfs -cat /apps/hive/warehouse/.snapshot/s20181204-164645.898/students/000000_0
```

Example output for a trivial table having two rows and three columns might look something like this:

```
fred flintstone351.28
barney rubble322.32
```

3. In Hive, insert the data into the table if the schema exists in the Hive warehouse; otherwise, restore the Hive Metastore, which includes the schemas, from the database dump you created in the pre-upgrade process.

## YARN Registry DNS instance fails to start

The YARN Registry DNS instance fails to start if another process on the host is bound to port 53. Ensure no other services that are binding to port 53 are on the host where the YARN Registry DNS instance is deployed.

### Steps

Class Loading issue when starting Solr. If you do not follow sequential steps during the upgrade, the Infra Solr instance might fail to start with an exception:

If you see the exception, follow the steps in this HCC article to work around the issue:

<https://community.cloudera.com/t5/Customer/ERROR-quot-null-org-apache-solr-common-SolrException-Error/tab/272751>

## Ambari Metrics System (AMS) does not start

When the Ambari Metrics System (AMS) does not start after upgrade, you can observe the following log snippet in the HBase Master:

```
master.HMaster: hbase:meta,,1.1588230740 is NOT online; state={1588230740 state=OPEN,
                                ts=1543610616273, server=regionserver1.domain.com,4121
                                3,1543389145213}; ServerCrashProcedures=true.
                                Master startup cannot progress, in holding-pattern until
                                region comes back online
```

The workaround is to manually clean up the Znode from ZooKeeper.

- If AMS mode = embedded, Remove the Znode data from local filesystem path, for example:

```
rm -f
/var/lib/ambari-metrics-collector/hbase-tmp/zookeeper/zookeeper_0/version-2/*
```

- If AMS mode = distributed, connect to the cluster ZooKeeper instance and delete the following node before restart:

```
/usr/hdp/current/zookeeper-client/bin/zkCli.sh -server localhost:2181  
[zk: localhost:2181(CONNECTED) 0] rmr /ams-hbase-unsecure/meta-region-server
```

## Ranger MySQL collation

Upgrading from HDP 2.6.5 to CDP 7.1.7 in some of the environments, the Ranger upgrade may fail. Such a failure can be due to MySQL collation used from the client-side or collation configuration enforced for all the clients from the server-side.

Error message during the upgrade failure: java.sql.SQLException: Illegal mix of collations (utf8\_general\_ci,IMPLICIT) and (utf8\_unicode\_ci,IMPLICIT) for operation '=' SQLException : SQL state: HY000 java.sql.SQLException: Illegal mix of collations (utf8\_general\_ci,IMPLICIT) and (utf8\_unicode\_ci,IMPLICIT) for operation '='  
ErrorCode: 1267

The MySQL collation issue occurs in those DB patches which contain dynamic SQL generation to perform the required DML operations. Some of the database patches that can fail in some environments are:

- 035-update-schema-for-x-policy.sql
- 036-denormalize-tag-tables.sql
- 047-sortorder-column-size.sql

Workaround: You must use the static SQL statements directly rather than generating them dynamically at the runtime.

Solution for DB patch 035-update-schema-for-x-policy.sql:

1. Open the 035-update-schema-for-x-policy.sql file in any editor.

2. From the SQL file, remove the code from line 65 to Line 110. For example, see the image below.

```
--
65 DROP PROCEDURE IF EXISTS removeConstraints;
66 DELIMITER ;;
67 CREATE PROCEDURE removeConstraints(vTableName varchar(128))
68 BEGIN
69     DECLARE done INT DEFAULT FALSE;
70     DECLARE cName VARCHAR(64);
71     DECLARE cur CURSOR FOR
72         SELECT DISTINCT CONSTRAINT_NAME
73         FROM INFORMATION_SCHEMA.Key_COLUMN_USAGE
74         WHERE TABLE_SCHEMA=DATABASE() AND TABLE_NAME = vTableName
75         AND REFERENCED_TABLE_NAME IS NOT NULL;
76     DECLARE CONTINUE HANDLER FOR NOT FOUND SET done = TRUE;
77
78     SET AUTOCOMMIT=0;
79     SET FOREIGN_KEY_CHECKS=0;
80
81     OPEN cur;
82
83     read_loop: LOOP
84         FETCH cur INTO cName;
85         IF done THEN
86             LEAVE read_loop;
87         END IF;
88         SET @sql = CONCAT('ALTER TABLE ',vTableName,' DROP FOREIGN KEY ',cName,');
89         PREPARE stmt FROM @sql;
90         EXECUTE stmt;
91         DEALLOCATE PREPARE stmt;
92     END LOOP;
93
94     CLOSE cur;
95
96     SET FOREIGN_KEY_CHECKS=1;
97     COMMIT;
98     SET AUTOCOMMIT=1;
99 END ;;
100 DELIMITER ;
101
102 call removeConstraints('x_policy_item');
103 call removeConstraints('x_policy_item_access');
104 call removeConstraints('x_policy_item_condition');
105 call removeConstraints('x_policy_item_datamask');
106 call removeConstraints('x_policy_item_group_perm');
107 call removeConstraints('x_policy_item_user_perm');
108 call removeConstraints('x_policy_item_rowfilter');
109 call removeConstraints('x_policy_resource');
110 call removeConstraints('x_policy_resource_map');
```

3. Copy the below code and paste the code from line number 65 in the SQL file: ALTER TABLE x\_policy\_item DROP FOREIGN KEY x\_policy\_item\_FK\_added\_by\_id; ALTER TABLE x\_policy\_item DROP FOREIGN KEY x\_policy\_item\_FK\_policy\_id; ALTER TABLE x\_policy\_item DROP FOREIGN KEY x\_policy\_item\_FK\_upd\_by\_id; ALTER TABLE x\_policy\_item\_access DROP FOREIGN KEY x\_policy\_item\_access\_FK\_added\_by\_id; ALTER TABLE x\_policy\_item\_access DROP FOREIGN KEY x\_policy\_item\_access\_FK\_atd\_id; ALTER TABLE x\_policy\_item\_access DROP FOREIGN KEY x\_policy\_item\_access\_FK\_pi\_id; ALTER TABLE x\_policy\_item\_access DROP FOREIGN KEY x\_policy\_item\_access\_FK\_upd\_by\_id; ALTER TABLE x\_policy\_item\_condition DROP FOREIGN KEY x\_policy\_item\_condition\_FK\_added\_by\_id; ALTER TABLE x\_policy\_item\_condition DROP FOREIGN KEY x\_policy\_item\_condition\_FK\_pcd\_id; ALTER TABLE x\_policy\_item\_condition DROP FOREIGN KEY x\_policy\_item\_condition\_FK\_pi\_id; ALTER TABLE x\_policy\_item\_condition DROP FOREIGN KEY x\_policy\_item\_condition\_FK\_upd\_by\_id; ALTER TABLE x\_policy\_item\_datamask DROP FOREIGN KEY x\_policy\_item\_datamask\_FK\_added\_by\_id; ALTER TABLE x\_policy\_item\_datamask DROP FOREIGN KEY x\_policy\_item\_datamask\_FK\_policy\_item\_id; ALTER TABLE x\_policy\_item\_datamask DROP FOREIGN KEY x\_policy\_item\_datamask\_FK\_type; ALTER TABLE x\_policy\_item\_datamask DROP FOREIGN KEY x\_policy\_item\_datamask



```

k_FK_upd_by_id; ALTER TABLE x_policy_item_group_perm DROP FOREIGN KEY x_policy_item_grou
p_perm_FK_added_by_id; ALTER TABLE x_policy_item_group_perm DROP FOREIGN KEY x_polic
y_item_group_perm_FK_group_id; ALTER TABLE x_policy_item_group_perm DROP FOREIGN KEY x_p
olicy_item_group_perm_FK_pi_id; ALTER TABLE x_policy_item_group_perm DROP FOREIGN KEY
x_policy_item_group_perm_FK_upd_by_id; ALTER TABLE x_policy_item_rowfilter DROP FOREIG
N KEY x_policy_item_rowfilter_FK_added_by_id; ALTER TABLE x_policy_item_rowfilter DROP FOREI
GN KEY x_policy_item_rowfilter_FK_policy_item_id; ALTER TABLE x_policy_item_rowfilter DR
OP FOREIGN KEY x_policy_item_rowfilter_FK_upd_by_id; ALTER TABLE x_policy_item_user_perm
DROP FOREIGN KEY x_policy_item_user_perm_FK_added_by_id; ALTER TABLE x_policy_item_
user_perm DROP FOREIGN KEY x_policy_item_user_perm_FK_pi_id; ALTER TABLE x_policy_item_u
ser_perm DROP FOREIGN KEY x_policy_item_user_perm_FK_upd_by_id; ALTER TABLE x_policy_ite
m_user_perm DROP FOREIGN KEY x_policy_item_user_perm_FK_user_id; ALTER TABLE x_polic
y_resource DROP FOREIGN KEY x_policy_resource_FK_added_by_id; ALTER TABLE x_policy_resou
rce DROP FOREIGN KEY x_policy_resource_FK_policy_id; ALTER TABLE x_policy_resource DROP
FOREIGN KEY x_policy_resource_FK_res_def_id; ALTER TABLE x_policy_resource DROP FOREIGN
KEY x_policy_resource_FK_upd_by_id; ALTER TABLE x_policy_resource_map DROP FOREIGN KEY
x_policy_resource_map_FK_added_by_id; ALTER TABLE x_policy_resource_map DROP FOREIGN KEY
x_policy_resource_map_FK_resource_id; ALTER TABLE x_policy_resource_map DROP FOREIGN KEY
x_policy_resource_map_FK_upd_by_id;

```

4. Save and close the file.

Solution for DB patch 036-denormalize-tag-tables.sql:

1. Open the 036-denormalize-tag-tables.sql file in any editor.

- From the SQL file, remove the code from line 40 to Line 82. For example, see the image

```

40 DROP PROCEDURE IF EXISTS removeConstraints;
41 DELIMITER ;;
42 CREATE PROCEDURE removeConstraints(vTableName varchar(128))
43 BEGIN
44     DECLARE done INT DEFAULT FALSE;
45     DECLARE cName VARCHAR(64);
46     DECLARE cur CURSOR FOR
47         SELECT DISTINCT CONSTRAINT_NAME
48         FROM INFORMATION_SCHEMA.Key_COLUMN_USAGE
49         WHERE TABLE_SCHEMA=DATABASE() AND TABLE_NAME = vTableName
50         AND REFERENCED_TABLE_NAME IS NOT NULL;
51     DECLARE CONTINUE HANDLER FOR NOT FOUND SET done = TRUE;
52
53     SET AUTOCOMMIT=0;
54     SET FOREIGN_KEY_CHECKS=0;
55
56     OPEN cur;
57
58     read_loop: LOOP
59         FETCH cur INTO cName;
60         IF done THEN
61             LEAVE read_loop;
62         END IF;
63         SET @sql = CONCAT('ALTER TABLE ',vTableName, ' DROP FOREIGN KEY ',cName, ';');
64         PREPARE stmt FROM @sql;
65         EXECUTE stmt;
66         DEALLOCATE PREPARE stmt;
67     END LOOP;
68
69     CLOSE cur;
70
71     SET FOREIGN_KEY_CHECKS=1;
72     COMMIT;
73     SET AUTOCOMMIT=1;
74 END ;;
75 DELIMITER ;
76
77 call removeConstraints('x_tag_attr_def');
78 call removeConstraints('x_tag_attr');
79 call removeConstraints('x_service_resource_element');
80 call removeConstraints('x_service_resource_element_val');
81
82 DROP PROCEDURE removeConstraints;|

```

- Copy the below code and paste the code from line number 40 in the SQL file: ALTER TABLE x\_service\_resource\_element DROP FOREIGN KEY x\_srvc\_res\_el\_FK\_added\_by\_id; ALTER TABLE x\_service\_resource\_element DROP FOREIGN KEY x\_srvc\_res\_el\_FK\_res\_def\_id; ALTER TABLE x\_service\_resource\_element DROP FOREIGN KEY x\_srvc\_res\_el\_FK\_res\_id; ALTER TABLE x\_service\_resource\_element\_val DROP FOREIGN KEY x\_srvc\_res\_el\_val\_FK\_add\_by\_id; ALTER TABLE x\_service\_resource\_element\_val DROP FOREIGN KEY x\_srvc\_res\_el\_val\_FK\_res\_el\_id; ALTER TABLE x\_service\_resource\_element\_val DROP FOREIGN KEY x\_srvc\_res\_el\_val\_FK\_upd\_by\_id; ALTER TABLE x\_tag\_attr DROP FOREIGN KEY x\_tag\_attr\_FK\_added\_by\_id; ALTER TABLE x\_tag\_attr DROP FOREIGN KEY x\_tag\_attr\_FK\_tag\_id; ALTER TABLE x\_tag\_attr DROP FOREIGN KEY x\_tag\_attr\_FK\_upd\_by\_id; ALTER TABLE x\_tag\_attr\_def DROP FOREIGN KEY x\_tag\_attr\_def\_FK\_added\_by\_id; ALTER TABLE x\_tag\_attr\_def DROP FOREIGN KEY x\_tag\_attr\_def\_FK\_tag\_def\_id; ALTER TABLE x\_tag\_attr\_def DROP FOREIGN KEY x\_tag\_attr\_def\_FK\_upd\_by\_id;
- Save and close the file.

Solution for DB patch 047-sortorder-column-size.sql:

- Open the 047-sortorder-column-size.sql file in any editor



```
policy_resource_map MODIFY COLUMN `sort_order` INT DEFAULT 0; ALTER TABLE x_resource_def
MODIFY COLUMN `sort_order` INT DEFAULT 0; ALTER TABLE x_resource_def MODIFY COLUMN
`sort_order` INT DEFAULT 0; ALTER TABLE x_resource_def MODIFY COLUMN `sort_order` INT
DEFAULT 0; ALTER TABLE x_resource_def MODIFY COLUMN `sort_order` INT DEFAULT 0; ALTER
TABLE x_resource_def MODIFY COLUMN `sort_order` INT DEFAULT 0; ALTER TABLE x_service_con
fig_def MODIFY COLUMN `sort_order` INT DEFAULT 0; ALTER TABLE x_service_config_def MODIF
Y COLUMN `sort_order` INT DEFAULT 0; ALTER TABLE x_service_config_def MODIFY COLUMN
`sort_order` INT DEFAULT 0; ALTER TABLE x_service_config_def MODIFY COLUMN `sort_order`
INT DEFAULT 0;
```

4. Save and close the file.

What next: You must continue the upgrade process.

## Rollback Ambari to 2.6.5

You must use the backup of Ambari 2.6.x to perform the following steps.

### About this task

#### Procedure

1. Stop all services in Ambari UI.
2. On the Ambari Server host, stop the Ambari Server.

```
ambari-server stop
Stop Ambari Agent on all host
ambari-agent stop
```

3. Obtain Ambari Repositories:

```
wget -nv
https://archive.cloudera.com/p/ambari/2.x/${your_ambari_version}/centos7/a
mbari.repo -O
/etc/yum.repos.d/ambari.repo
yum clean all
```

4. Downgrade ambari-server, ambari-agent packages manually using “yum downgrade <packagename>”

```
yum downgrade ambari-server
yum downgrade ambari-agent # on all hosts
```

5. Reinitialize Ambari DB - Drop the old schema and initialize blank new schema :

Example - Postgres

```
psql -W -d ${AMBARI_DB_NAME} --user ${AMBARI_DB_USER}
DROP SCHEMA ${AMBARI_DB_NAME} CASCADE;
CREATE SCHEMA ambari;
GRANT ALL ON SCHEMA ambari TO ambari;
```

MYSQL

```
mysql -u ${AMBARI_DB_USER}
drop database ${AMBARI_DB_NAME};
create database ${AMBARI_DB_NAME};
```

- Populate the Ambari database with backup data:

POSTGRES

```
psql -W -d ${AMBARI_DB_NAME} --user ${AMBARI_DB_USER} < pre_upgrade.sql
```

MYSQL

```
mysql -u ${AMBARI_DB_USER} ${AMBARI_DB_NAME} < pre_upgrade.sql
```

- Clear views cache (it has views for 7.1 cached, some of which are not supported):

```
rm -f /var/lib/ambari-server/resources/views/work/*
```

- Restore these files from the backup:

```
/etc/ambari-server/conf/ambari.properties
/etc/ambari-agent/conf/ambari-agent.ini
```

- Restore old stack symlinks on Agent nodes:

```
hdp-select set all 2.6.5.0-292
```

- Check ambari-server user:

```
grep ambari-server.user /etc/ambari-server/conf/ambari.properties
```

If ambari-server is running as non-root user, set permissions for files by running: ambari-server setup

- Start Ambari

## Rollback HDP Services

After Ambari rollback, you must rollback the HDP services in the following order to restore to HDP 2.6.5.

### Overview

You can roll back an upgrade from HDP 2.6.5.x to CDP Private Cloud Base 7.1.6, 7.1.7 SP2, and 7.1.8. The rollback restores your HDP cluster to the state it was in before the upgrade, including Kerberos and TLS/SSL configurations.

Before you start rolling back the CDP Private Cloud Base 7 to HDP 2, review the following information.

Caveats

- Any data created after the upgrade is lost. In a typical upgrade, you first upgrade Ambari from version 2.6.5 to version 7.1.6, and then you use the upgraded version of Ambari 7.1.6 to upgrade HDP 2 to CDP Private Cloud Base 7. (See [Upgrading a Cluster](#)) If you want to roll back this upgrade, follow these steps to roll back your cluster to its state prior to the upgrade.
- Follow all of the steps in the order presented in this topic. Cloudera recommends that you read through the backup and rollback steps before starting the backup process. You may want to create a detailed plan to help you anticipate potential problems.
- You can roll back to HDP 2 after upgrading to CDP Private Cloud Base 7 only if the HDFS upgrade has not been finalized. The rollback restores your HDP cluster to the state it was in before the upgrade, including Kerberos and TLS/SSL configurations.
- These rollback steps depend on complete backups taken before upgrading Ambari and HDP. See [Back Up Ambari](#) and [Backing up HDP cluster](#). For steps where you need to restore the contents of a directory, clear the contents of the directory before copying the backed-up files to the directory. If you fail to do this, artifacts from the original upgrade can cause problems if you attempt the upgrade again after the rollback.

Review Limitations

The rollback procedure has the following limitations. This rollback guide covers rolling back upgrade before finalizing the HDP7 upgrade in Ambari and does not cover the AM2CM Stage.

- HDFS – If you have finalized the HDFS upgrade, you cannot roll back your cluster.
- Configuration changes, including the addition of new services or roles after the upgrade are not retained after rolling back Ambari. Cloudera recommends that you not make configuration changes or add new services and roles until you have finalized the HDFS upgrade and no longer require the option to roll back your upgrade.
- HBase – If your cluster is configured to use HBase replication, data written to HBase after the upgrade might not be replicated to peers when you start your rollback. This topic does not describe how to determine which, if any, peers have the replicated data and how to roll back that data. For more information about HBase replication, see HBase Replication.
- Kafka – Once the Kafka log format and protocol version configurations (the `inter.broker.protocol.version` and `log.message.format.version` properties) are set to the new version (or left blank, which means to use the latest version), Kafka rollback is not possible.

Unsupported components

For more information on supported components, see [unsupported components](#).

## ZooKeeper

Use the following step to restore the backed-up data for ZooKeeper. For zookeeper, the default folder is `/hadoop/zookeeper`.

### Procedure

1. Execute the following command to clean up and restore:

```
rm -rf /hadoop/zookeeper
tar xf up.tar.gz -C /
```

2. Start Zookeeper in Ambari UI.

## Ambari-Metrics

To restore the backed-up Ambari-Metrics data, downgrade the ambari-metrics packages on all hosts and restore the backed-up-ambari-metrics data.

### Procedure

1. Downgrade ambari-metrics packages on all hosts where they are installed. On every host in your cluster running a Metrics Monitor, run the following commands:

```
yum downgrade ambari-metrics-monitor ambari-metrics-hadoop-sink
```

Execute the following command on all hosts running the Metrics Collector:

```
yum downgrade ambari-metrics-collector
```

Execute the following command on the host running the Grafana component:

```
yum downgrade ambari-metrics-grafana
```

2. Restore backed-up ambari-metrics data on the host where the ambari-metrics-collector is installed

```
rm -rf /var/lib/ambari-metrics-collector
tar xf ams-backup.tar.gz -C /
```

3. Start Ambari-Metics in Ambari UI.

## Ambari Infra Solr

Use the following steps to restore backed-up ambari infrasolr data.

### Procedure

1. Downgrade ambari-infra-solr, ambari-infra-solr-client packages on all hosts where they are installed.

Execute the following command on all hosts running the infra Solr Instance:

```
yum downgrade ambari-infra-solr
```

Execute the following command on all hosts running the infra Solr Client:

```
yum downgrade ambari-infra-solr-client
```

2. Restore backed-up infra-solr data on the node where the infra-solr is installed.

Start Ambari-Infra service in Ambari UI then execute the following commands:

For unsecured cluster

```
curl -v
"http://${INFRA_SOLR_URL}/solr/vertex_index/replication?command=restore
&location=/path/to/backup/directory&name=vertex_index_backup"

curl -v
"http://${INFRA_SOLR_URL}/solr/edge_index/replication?command=restore&lo
cation=/path/to/backup/directory&name=edge_index_backup"

curl -v
"http://${INFRA_SOLR_URL}/solr/fulltext_index/replication?command=restore
&location=/path/to/backup/directory&name=fulltext_index_backup"
curl -v
"http://${INFRA_SOLR_URL}/solr/ranger_audits/replication?command=restor
e&location=/path/to/backup/directory&name=ranger_audits_backup"
curl -v
"http://${INFRA_SOLR_URL}/solr/hadoop_logs/replication?command=restore&
location=/path/to/backup/directory&name=hadoop_logs_backup"
curl -v
"http://${INFRA_SOLR_URL}/solr/audit_logs/replication?command=restore&loca
tion=/path/to/backup/directory&name=audit_logs_backup"

curl -v
"http://${INFRA_SOLR_URL}/solr/history/replication?command=restore&loca
tion=/path/to/backup/directory&name=history_backup"
```

For secured cluster

If the cluster is Kerberized, then you must kinit as the service principal.

```
curl -v --negotiate -u:
"http://${INFRA_SOLR_URL}/solr/vertex_index/replication?command=restore
&location=/path/to/backup/directory&name=vertex_index_backup"

curl -v --negotiate -u:
"http://${INFRA_SOLR_URL}/solr/edge_index/replication?command=restore&lo
cation=/path/to/backup/directory&name=edge_index_backup"

curl -v --negotiate -u:
"http://${INFRA_SOLR_URL}/solr/fulltext_index/replication?command=restore
&location=/path/to/backup/directory&name=fulltext_index_backup"
curl -v --negotiate -u:
```

```
"http://${INFRA_SOLR_URL}/solr/ranger_audits/replication?command=restore&location=/path/to/backup/directory&name=ranger_audits_backup"

curl -v --negotiate -u:
"http://${INFRA_SOLR_URL}/solr/hadoop_logs/replication?command=restore&location=/path/to/backup/directory&name=hadoop_logs_backup"

curl -v --negotiate -u:
"http://${INFRA_SOLR_URL}/solr/audit_logs/replication?command=restore&location=/path/to/backup/directory&name=audit_logs_backup"

curl -v --negotiate -u:
"http://${INFRA_SOLR_URL}/solr/history/replication?command=restore&location=/path/to/backup/directory&name=history_backup"
```

3. Start Ambari-Infra in Ambari UI.

## Ranger

Rollback procedure of Ranger includes restoring Ranger admin database and Ranger KMS database.

### Restore Ranger Admin Database

To restore Ranger admin databases, perform the following steps. First Stop Ranger Admin and KMS services if it's running and then follow below given steps.

#### MySQL

Perform the following steps to restore a database.

1. Delete an existing Database.
2. Create an empty new database on the database host.
3. Restore the database using below mysql command.

```
mysql -u root
drop database ranger;
create database ranger;
GRANT ALL PRIVILEGES ON ranger.* TO 'rangeradmin'@'localhost';
$ mysql -u [username] -p existing_empty_db_name < dump_fileName.sql
```

#### Example

```
mysql -u rangeradmin -p rangeradmin < /root/backups/ranger/db/admin/ranger.sql
```

Select the Enter key. Type the database password at the password prompt.

#### POSTGRES

Perform the following steps to restore a database.

1. Delete an existing Database.
2. Create an empty new database in its place.
3. Execute the below command on postgres db host.

```
dropdb -U owner_username dbname; [Enter db owner password at the prompt]
```

#### Example

```
dropdb -U rangeradmin ranger;
```

```
su - postgres
psql
```



```
create database ranger;
ALTER DATABASE ranger OWNER TO rangeradmin;
\q
exit
psql -U rangeradmin ranger < /root/backups/ranger/db/admin/ranger.sql
```

[Enter db owner password at the prompt as rangeradmin]

Set path to oracle home if required :

```
export ORACLE_HOME=/opt/oracle/product/12.2.0
export PATH=${PATH}:${ORACLE_HOME}/bin
export ORACLE_SID=orcl12c
```

Restore ranger admin database.

```
rm -rf del_admin_tbl_cmd.sql
sqlplus -s rangeradmin/rangeradmin << EOF
  spool on
  spool del_admin_tbl_cmd.sql
  SELECT 'DROP TABLE "' || TABLE_NAME || '" CASCADE CONSTRAINTS;'
FROM user_tables
  union ALL
  select 'drop ' || object_type || ' ' || object_name || ';' from us
er_objects
  where object_type in ('VIEW','PACKAGE','SEQUENCE', 'PROCEDURE',
'FUNCTION')
  union ALL
  SELECT 'drop '
  || object_type
  || ' '
  || object_name
  || ' force;'
FROM user_objects
WHERE object_type IN ('TYPE');
  spool off
  @del_admin_tbl_cmd.sql
  exit;
                                EOF
```

Type the database password at the password prompt and then execute the below imp command.

```
imp rangeradmin/rangeradmin file=backups/ranger/db/admin/orcl12c.sql log
=backups/ranger/db/admin/restore.log
```

### Restore Ranger KMS Database

Restoring Ranger KMS involves steps for restoring MSQl, POSTGRES, and Oracle databases.

#### MySQL

To restore a database do the following:

- Delete an existing database.
- Create an empty new database on the Database host.
- Restore the database using below msyql command.

```
mysql -u root
drop database rangerkms;
create database rangerkms;
GRANT ALL PRIVILEGES ON rangerkms.* TO 'rangerkms'@'localhost';
$ mysql -u [username] -p existing_empty_db_name < dump_fileName.sql
```

## Example

```
mysql -u rangerkms -p rangerkms < /root/backups/ranger/db/kms/rangerkms.sql
```

Select the Enter key. Type the database password at the password prompt.

## POSTGRES

To restore data, do the following:

- Delete an existing database.
- Create an empty new database in its place.
- Execute the below command on postgres database host.

```
dropdb -U owner_username dbname; [Enter db owner password at the prompt]
```

## Example

```
dropdb -U rangerkms rangerkms;
su - postgres
psql
create database rangerkms;
ALTER DATABASE rangerkms OWNER TO rangerkms;
\q
exit
psql -U rangerkms rangerkms < /root/backups/ranger/db/kms/rangerkms.sql
[Enter db owner password at the prompt as rangeradmin]
```

## For oracle

```
rm -rf del_kms_tbl_cmd.sql
sqlplus -s rangerkms/rangerkms << EOF
    spool on
    spool del_kms_tbl_cmd.sql
    SELECT 'DROP TABLE "' || TABLE_NAME || '" CASCADE CONSTRAINTS;' FR
OM user_tables
        union ALL
        select 'drop ' || object_type || ' ' || object_name || ';' from user_o
bjects
        where object_type in ('VIEW','PACKAGE','SEQUENCE', 'PROCEDURE', '
FUNCTION')
        union ALL
        SELECT 'drop '
        || object_type
        || ' '
        || object_name
        || ' force;'
        FROM user_objects
        WHERE object_type IN ('TYPE');
    spool off
@del_kms_tbl_cmd.sql
exit;
EOF
```

Select Enter and then execute the below imp command.

```
imp rangerkms/rangerkms file=backups/ranger/db/kms/orcl12c.sql
log=backups/ranger/db/kms/restore.log
```

After rollback, follow the below steps to add ranger LB members node principals in spnego keytab.

Execute on LB node (On ranger external url host) :

```
kadmin.local
ktadd -norandkey -kt /etc/security/keytabs/spnego.service.keytab HTTP/admin-
host1@EXAMPLE.COM
ktadd -norandkey -kt /etc/security/keytabs/spnego.service.keytab HTTP/admin-
host2@EXAMPLE.COM
```

Copy same keytab on ranger-admin nodes from LB host:

```
scp /etc/security/keytabs root@admin-host1:/etc/security/keytabs/spnego.ser
vice.keytab
scp /etc/security/keytabs root@admin-host2:/etc/security/keytabs/spnego.serv
ice.keytab
```

If you've performed Ambari-Infra rollback already, there are no additional rollback steps required to restore the Solr Collection.

Start Ranger and RangerKMS in Ambari UI.

## HDFS

Before starting the rollback procedure, make sure that all HDFS service roles are stopped.

### About this task



**Note:** Before the HDFS rollback, Zookeeper, Ranger, Ambari-Metrics, Ambari-Infra has to be rolled back and started in Ambari UI.

### Procedure

1. Roll back all the JournalNodes. (Only required for clusters where high availability is enabled for HDFS). Use the JournalNode backup that you have created when you backed up HDFS before upgrading to the CDP Private Cloud Base.
  - a) Log in to each JournalNode host and do the following:
    1. remove the `$(dfs.journalnode.edits.dir)/current` directory
    2. restore the backup of `$(dfs.journalnode.edits.dir)/current` into `$(dfs.journalnode.edits.dir)/current` into `$(dfs.journalnode.edits.dir)/current`
2. Note down the target of the `/etc/hadoop/confsymbolic` link and remove it
3. Move the backup of `/etc/hadoop/conf` back to its original place, and do these steps on all the cluster nodes where HDFS roles are installed, so on all NameNodes, JournalNodes and DataNodes.

#### 4. Roll back all of the NameNodes.



**Note:** if high availability is not enabled on your cluster, then leave the Secondary NameNode as it is for now.

Use the backup of the Hadoop config directory you created during the backup phase.

Perform the following steps on all NameNode hosts:

- a) Start FailoverControllers, and JournalNodes
  - b) If you use Kerberos authentication, authenticate with kinit with the NameNode's principal, otherwise change to the hdfs service user (usually `sudo -u hdfs`)
  - c) Run the following command: `hdfs namenode -rollback`
  - d) Restart HDFS FailoverControllers, and JournalNodes in Ambari, then start the NameNodes note that one of the NameNodes should start, and one of them will remain in starting state. When one of the NameNodes are marked as started proceed to DataNode rollback.
5. Roll back all of the DataNodes. Use the backup of the Hadoop config directory you created during the backup phase. Perform the following steps on all DataNode hosts:

- a) If you use Kerberos authentication, authenticate with kinit with the NameNode's principal, otherwise change to the hdfs service user (usually `sudo -u hdfs`)
- b) Run the following commands:
  - `export HADOOP_SECURE_DN_USER=<hdfs service user>`
  - `export JSVC_HOME=<directory of jsvc binary>` usually the directory is `/usr/lib/bigtop-utils/`
  - `hdfs datanode -rollback`
  - Look for output from the command similar to the following that indicates when the DataNode rollback is complete. wait until all storage directories are rolled back:

```
INFO common.Storage: Layout version rolled back to -56 for storage /
storage/dir_x
INFO common.Storage (DataStorage.java:doRollback(952)) - Rollback
of /storage/dir_x is complete
```



**Note:** If you do not see the output, check for the `jsvc.err` file in the DataNode's log directory. If you see these log messages in the output, or in the `jsvc.err` file for all of your DataNode data folders, then stop the process by typing `ctrl+c` as the DataNode rollback is ready.

6. If your cluster is not configured for NameNode High Availability, roll back the Secondary NameNode. Perform the following steps on the Secondary NameNode host:
  - a) Move the Secondary NameNode data directory to a backup location. (`$(dfs.namenode.name.dir)`)
  - b) If you use Kerberos authentication, authenticate with kinit with the NameNode's principal, otherwise change to the hdfs service user (usually `sudo -u hdfs`)
  - c) Run the following command: `hdfs secondarynamenode -format`

After rolling back the Secondary NameNode, terminate the console session by typing Control-C. Look for output from the command similar to the following that indicates when the DataNode rollback is complete:

```
INFO namenode.SecondaryNameNode: Web server init done
```

7. Restore the original symlink with the noted target as `/etc/hadoop/conf` on all the nodes where it has changed.
8. Restart the HDFS service. Open Ambari, and go to the HDFS service page, in the Service actions dropdown select Start.
9. Monitor the service, and if everything comes up fine, check the HDFS file system availability, you can run an `hdfs fsck /` or generate the file system listing with `hdfs dfs -ls -R /` and compare it with the one that you did as part of the backup procedure to see if everything got rolled back properly. In case of any issues, please contact Cloudera Support before you proceed.

## YARN

If you have performed HDFS rollback already, there are no additional rollback steps required for YARN.

Start Yarn in Ambari UI

## HBase

If you have performed HDFS rollback already, there are no additional rollback steps required for HBase.

Start HBase in Ambari UI.

## Kafka

To roll back Kafka, perform the following steps.

### Procedure

1. Kafka service depends on Zookeeper. Make sure Zookeeper data is restored.
2. After rollback, start the Kafka service, and check, if the producers and consumers can connect to the cluster.
3. Remove the inter broker protocol and log format version settings from the Kafka settings:
  - a) Log in to Ambari.
  - b) Choose the Kafka service.
  - c) Select the Configs page.
  - d) Find the Custom kafka-broker section.
  - e) Remove following properties:
    - `inter.broker.protocol.version=current_Kafka_version`
    - `log.message.format.version=current_Kafka_version`
  - f) Restart the Kafka service.
  - g) Start Kafka in Ambari UI

## Atlas

Perform the following steps to restore HBase tables and ATLAS\_ENTITY\_AUDIT\_EVENTS table.

### Procedure

1. Stop Atlas from Ambari
2. Setup Kerberos Credentials in case of secured cluster . Locate the atlas user's keytab and use kinit to cache the kerberos ticket for atlas user .

Example

```
kinit -kt path/to/atlas.service.keytab atlas/hostname@domain
```

3. Restore the HBase tables, atlas\_titan and ATLAS\_ENTITY\_AUDIT\_EVENTS.
4. If you've performed Ambari-Infra rollback already, there are no additional rollback steps required to restore the Solr Collections.
5. Start Atlas in Ambari UI.

### Restore HBase Tables

Do the following to restore atlas\_titan table.

```
> disable 'atlas_titan'  
> restore_snapshot 'atlas-titan-backup'  
> enable 'atlas_titan'
```

### Restore ATLAS\_ENTITY\_AUDIT\_EVENTS table

Do the following to restore ATLAS\_ENTITY\_AUDIT\_EVENTS table.

```
> disable 'ATLAS_ENTITY_AUDIT_EVENTS'  
> restore_snapshot 'ATLAS_ENTITY_AUDIT_EVENTS-backup'  
> enable 'ATLAS_ENTITY_AUDIT_EVENTS'
```

## Hive

You must delete your existing Hive Metastore database to roll back the Hive services.

To restore data on the node where the Hive Metastore database is located, perform the following steps.

1. Delete an existing database.
2. Create an empty database in its place.

```
$ mysql -u <hive_user>  
drop database <hive_db>;  
create database <hive_db>;  
$ mysql -u <hive_user> <hive_db> </path/to/dump_file>
```

If you've performed HDFS rollback already, there are no additional rollback steps required for Hive.

Start Hive service in Ambari UI.

## Spark

Know more about the Spark roll back.

- The Spark application history lives in HDFS, and with the rollback of HDFS, the history at the time of the backup is restored. Any Spark applications run after the backup and rollback will not be visible.
- After the rollback, versions of Spark applications built against HDP 2 should be used instead of versions that are rebuilt against CDP.
- Start Spark in Ambari UI

## Oozie

To roll back the Oozie service, you must restore the Oozie database.

- Start Oozie in Ambari UI.

It is needed to re-generate the oozie.ha.keytab file using the spnego.service.keytab files from the Oozie server hosts and from the Oozie load balancer host and then distribute the generated oozie.ha.keytab onto the respective Oozie hosts. In order to achieve this, follow the steps listed in this [website](#).

- Restart Oozie.

## Knox

With the backup and rollback of Ambari, Knox is also backed up and rolled back by default.

Start Knox in Ambari UI.

## Zeppelin

With the backup and rollback of HDFS, Zeppelin is also backed up and rolled back by default.

Start Zeppelin in Ambari UI.

## Log Search

### Procedure

1. On every host in your cluster running a Log Search Server, run the following commands:

```
yum downgrade ambari-logsearch-portal
```

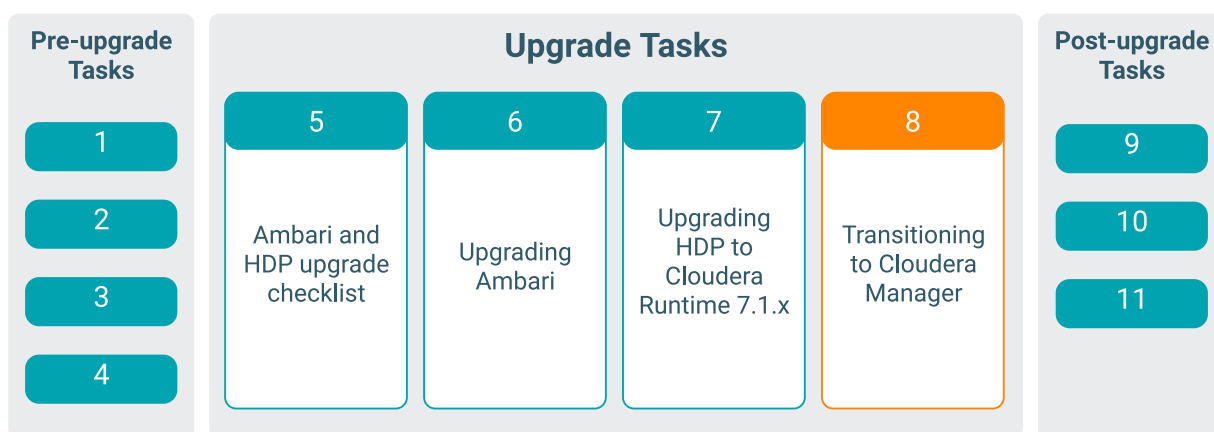
2. Execute the following command on all hosts running the Logfeeder:

```
yum downgrade ambari-logsearch-logfeeder
```

3. Start Log Search in Ambari UI.

## Transitioning to Cloudera Manager

This section helps you to transition from HDP 2.6.5.x to Cloudera Manager and Cloudera Private Cloud Base



The high level of steps to transition from the HDP 2.6.5 cluster to Cloudera Manager:

## Transitioning to Cloudera Manager. Transition from HDP 2 to Cloudera Manager.

- a [Perform the pre-transition steps for the components deployed in your clusters including exporting blueprint with hosts from Ambari.](#)
- b [Install Cloudera Manager 7.x.x - the Cloudera Manager server, daemons, and agents.](#)
- c [Transitioning the HDP cluster to CDP Private Cloud Base cluster using the AM2CM tool.](#)
- d [Start CDP Private Cloud Base using Cloudera Manager.](#)
- e [Perform the post transition steps in Cloudera Manager.](#)

### Pre-transition steps

From Ambari configurations, collect passwords of databases, keystores, and IDM services for components that are not part of the transition. This information will help in starting the services in Cloudera Manager after transition.



**Note:** The AM2CM tool does not migrate the Ambari Infra service to Cloudera Manager. However, the tool adds a Solr service in place of Ambari Infra to the Cloudera Manager-managed CDP cluster. The tool does not migrate Ambari Infra configurations and data to Solr in Cloudera Manager.



#### Important:

- Backup custom files from /etc/hadoop/ in the HDP cluster. For example, /etc/hadoop/conf/topology\_script.py
- Ensure no components or hosts are in decommissioned state.
- Ambari Metrics will not be migrated from Ambari to Cloudera Manager.
- If are using the AM2CM tool <http://archive.cloudera.com/am2cm/1.x/am2cm-1.0.0.0-58.tar.gz>, then this AM2CM tool supports:
  - HDFS (HA), YARN (HA), HIVE (HA), Zookeeper (HA), HBase (HA), Oozie (HA)
  - Spark, MR2, Tez, Sqoop, Livy, Kafka, CM
- The AM2CM 2.0.0.0 tool supports DellEMC PowerScale ISILON migration.



**Important:** For log4j configuration migrations, AM2CM tool migrates only log file names, location, backup index , size, and logging level. All other details will not be migrated or translated.



**Knox SSO on Ambari and Cloudera Manager:** After you upgrade your cluster with Knox SSO and proxy enabled, the post upgrade Cloudera Manager tests fail because it is going through Knox SSO. The workaround is to disable the Ambari SSO before migrating the cluster because Ambari- specific configurations are not migrated to Cloudera Manager. You must explicitly enable SSO in Cloudera Manager after migration. For more information, see [Proxy Cloudera Manager through Apache Knox](#).

## Databases

Before you start upgrading from Ambari HDP to Cloudera Manager Cloudera Private Cloud Base, you must ensure that you have a backup of the following component databases: Hive Metastore, Ranger, and Oozie.

Cloudera recommends that you transfer data from the old databases to a new set of databases, and use the new databases as reference in Cloudera Private Cloud Base. You can retain the old databases for reference, and to research any issues that you may encounter when running your applications in Cloudera Private Cloud Base.

## Kerberos

Prior to installing Kerberos on Cloudera Manager, you must collect the details from the Kerberos installation in Ambari.

Collect the values for the following settings:

- KDC Type (AD, MIT KDC, Redhat IPA)
- Admin Account Username
- Admin Account Password
- Kerberos Domain
- Kerberos Realm
- KDC Host
- KDC Admin Host
- Service Port (if non standard)
- Encryption Type (default 'rc4-hmac')
- AD Directory Suffix

Cloudera Manager maintains a single principal for all the roles in a service. Hence Auth\_to\_local rules are not migrated.

Service user names in Ambari will be migrated to Kerberos principal names in Cloudera manager as part of the transition process.

### Kerberos principal

YARN service and Kafka service in Cloudera Manager does not take non-default principal names. This section is not applicable if you are upgrading to CDP Private Cloud Base 7.1.7.

Currently, the YARN service in Cloudera Manager does not take non-default principal names. The YARN principle name is not migrated to Cloudera Manager if you use any custom principal name in the YARN service on Ambari.

Currently, the Kafka service in Cloudera Manager does not take non-default principal names. The Kafka principle name is not migrated to Cloudera Manager if you use any custom principal name in the Kafka service on Ambari.

## Atlas migration

Perform the pre-transition steps for Atlas.

### About this task

When you are upgrading from HDP 2.6.x to HDP intermediate bits, ensure you set the atlas.migration.data.filename=<non-existent-path> value for the atlas-application.properties configuration For example, /tmp/nonexistentpath. You must set the correct path for <the actual exported migration data path> before you upgrade to Cloudera Manager managed CDP Private Cloud Base cluster.

### Procedure

1. In Ambari, select Services > Atlas > Configs > Advanced > Custom application-properties.
2. Click Add Property, and add the atlas.migration.data.filename property.

### Example

atlas.migration.data.filename=</tmp/non-existent-path>

## HDFS

Perform the following pre-transition steps.

Do not change the HA namespace during the upgrade. Changing HA namespace impacts Hive metadata locations and possibly HBase (if the location reference includes the full HDFS location).

### Preparing HDFS

Make sure HDFS is healthy.

### Before you begin

To ensure HDFS is healthy

### Procedure

- Save the Namespace
- Ensure that the under-replicated blocks do not exist
- Ensure that the Exclude hosts list is empty. Before upgrading, remove the decommissioned datanodes from HDFS.

### Backup the non-default Rack Awareness Topology script

Backup any custom configurations of the rack awareness topology script.

In the Ambari HDFS service, if net.topology.script.file.name is configured with a non default topology script (that is not the default /etc/hadoop/conf/topology\_script.py script), then take a backup of it.

After transitioning to Cloudera Manager, you can reconfigure your script.



**Note:** The reason for this is that the default script has a different name in Cloudera Manager. That is the default script is the same in Ambari and Cloudera Manager but the two management systems have different names for that same script.

## Spark

Run the following commands on all the hosts where spark binaries are present on the filesystem.

Perform the following steps:

```
mkdir /usr/bin/sparkbackup
mv /usr/bin/spark-shell /usr/bin/sparkbackup/
mv /usr/bin/spark-submit /usr/bin/sparkbackup/
mv /usr/bin/pyspark /usr/bin/sparkbackup/
mv /usr/bin/spark-class /usr/bin/sparkbackup/
mv /usr/bin/spark-script-wrapper.sh /usr/bin/sparkbackup/
mv /usr/bin/spark-sql /usr/bin/sparkbackup/
```

### Spark2/Livy

Check the Livy server component.

Cloudera Manager does not allow more than one Livy server component. Remove if Ambari has more than one Livy server component before migration.

## Ranger

Configure Ranger in Cloudera Manager-managed Cloudera Private Cloud Base cluster.

Usually, Ranger and most HDP services are set up to synchronize and authenticate against an LDAP/AD service. In the Cloudera Manager-managed Cloudera Private Cloud Base cluster, additional options are available to manage integrations using the aspects of the existing operating system's IDM integration.

The configuration for Ranger in Cloudera Manager-managed Cloudera Private Cloud Base cluster is through PAM (Pluggable Authentication Module). In a Kerberized environment, PAM integration is available at the Operating system level.

You can use the Ranger services for authentication, user, and group synchronization using the PAM integration.

If PAM integration is not an option, then you must record the details from Ambari to install and configure Ranger in Cloudera Manager.

## Kerberos - Optional task

De-kerberize the Ambari HDP cluster and kerberize the Cloudera Manager cluster.

You can de-kerberize the Ambari HDP cluster before migration and kerberize the Cloudera Manager cluster when the Ambari HDP migration is complete.



**Note:** Even without kerberos, migration works.

## Solr

The AM2CM tool does not migrate Ambari Infra Solr configuration and data to Cloudera Manager. However, it adds a new Solr service. To take a backup, perform the following steps after the transition.

1. Create a local directory for backups first:

```
mkdir -p [backup_dir]
```

2. Execute the following commands on the hosts where Solr is installed.

```
mv /etc/alternatives/solr-conf [backup_dir]/solr-confhdp
mv /etc/alternatives/hue-conf [backup_dir]/hue-confhdp
mv /etc/alternatives/hbase-solr-conf [backup_dir]/hbase-solr-confhdp
```

For RHEL:

```
mv /var/lib/alternatives/solr-conf [backup_dir]/solr-confhdp
```

For Ubuntu:

```
mv /var/lib/dpkg/alternatives [backup_dir]/solr-confhdp
```

For Sles:

```
mv /var/lib/rpm/alternatives [backup_dir]/solr-confhdp
```

## Cloudera Manager Installation and Setup

Install Cloudera Manager, install Cloudera Manager agent and daemons, add Cloudera Management service, enable TLS, and finally configure clusters to use Kerberos, if you are transitioning from Ambari 7.1.x.x to Cloudera Manager 7.3.1 using the AM2CM tool.

## About this task



**Important:** If you are transitioning from Ambari 7.1.x.x to Cloudera Manager 7.2.4 using the AM2CM tool, then you must see, [Cloudera Manager 7.2.4](#) documentation.

## Procedure

1. Prepare to install and configure the Cloudera Manager packages. For more information, see [Configuring Repository](#). Do this if you have not done it already. Confirm that the repo is set up.
2. Install Cloudera Manager Server. For more information on installing Cloudera Manager Server, see [Installing Cloudera Manager](#).



**Note:** During the upgrade process, you can place Cloudera Manager and its related Cloudera Manager Services components on the same node as Ambari. Ensure that the node has sufficient capacity to temporarily run Cloudera Manager and Ambari in parallel. If you do not wish to colocate these services, you can place them on separate management nodes in the cluster.

3. Preconfigure the databases for:
  - Ranger
  - Cloudera Manager Server
  - Cloudera Management Service roles - Reports Manager
  - Data Analytics Studio (DAS) Supported with PostgreSQL only.
  - Hue
  - Each Hive metastore
  - Oozie
  - Schema Registry
  - Streams Messaging Manager

For more information, see [Setup Cloudera Manager database](#) and [Install and Configure Databases](#).

4. Install Agent on all hosts in the cluster. It is possible to add hosts to Cloudera Manager using the [Installation Wizard](#).
5. Start Cloudera Manager Server and Cloudera Manager agent on all hosts. For more information, see [Cloudera Manager Agent](#) and [Cloudera Manager Server](#).
6. Install Cloudera Manager User licence. For more information, see [Installation Wizard](#). (Upload the license file and

exit the cluster setup by clicking the Cloudera Manager icon )



### Caution:

- Do not set up a cluster using the Wizard (Step 7) .
  - Do not proceed to Welcome (Add Cluster - Installation).
7. Add Hosts to a cluster. To add hosts to a cluster, see [Adding Hosts to a cluster](#).
  8. Add Cloudera Manager management service to the cluster. To add services to the cluster, see [Select Services](#).
  9. Set up Kerberos. If you have a Kerberos cluster, then you must add the KDC details in the Administration>Security>Kerberos Credentials>Setup KDC for Cloudera Manager page using Cloudera Manager. For more information on Kerberos and Active Directory, see [Enabling Kerberos authentication for CDP](#)



### Note:

- You must set the value of Maximum Renewable Life for Principals to 0 or to the value that was provided for Ambari KDC. If Maximum Renewable Life for Principal and Ambari KDC do not match, keytabs generation fails in Cloudera Manager.
- If you are not upgrading from CDH or HDP to CDP Private Cloud Base, you can follow the [Production Installation](#) instructions to perform a fresh installation in your production environment.

### Related Information

[Configuring Repository](#)  
[Cloudera Manager Server](#)  
[Installation Wizard](#)  
[How to Configure Clusters to Use Kerberos for Authentication](#)  
[Software download matrix](#)

## Installing JDBC Driver

You must install the required JDBC driver.

Download, extract, and copy the JDBC driver, renamed, to `/usr/share/java/`. If the target directory does not yet exist, create it.

### Installing the Postgres JDBC Driver

1. Install the PostgreSQL JDBC driver. If you would like to use the PostgreSQL JDBC driver version shipped with the OS repositories, run the following command:

```
yum install postgresql-jdbc*
```

You can also download the JDBC driver from the official PostgreSQL JDBC Driver website – <https://jdbc.postgresql.org>.

2. Rename the Postgres JDBC driver `.jar` file to `postgresql-connector-java.jar` and copy it to the `/usr/share/java` directory. The following copy command can be used if the Postgres JDBC driver `.jar` file is installed from the OS repositories:

```
cp /usr/share/java/postgresql-jdbc.jar /usr/share/java/postgresql-connector-java.jar
```

3. Confirm that the `.jar` file is in the Java share directory:

```
ls /usr/share/java/postgresql-connector-java.jar
```

4. Change the access mode of the `.jar` file to 644:

```
chmod 644 /usr/share/java/postgresql-connector-java.jar
```

### Installing the MySQL JDBC Driver for MariaDB

The MariaDB JDBC driver is not supported. Follow the steps in this section to install and use the MySQL JDBC driver instead.

Install the JDBC driver on the Cloudera Manager Server host, as well as any other hosts running services that require database access.




**Note:** If you already have the JDBC driver installed on the hosts that need it, you can skip this section. However, MySQL 5.6 requires a version 5.1.26 or higher.

Cloudera recommends that you consolidate all roles that require databases on a limited number of hosts, and install the driver on those hosts. Locating all such roles on the same hosts is recommended but not required. Make sure to install the JDBC driver on each host running roles that access the database.



**Note:** Cloudera recommends using only version 5.1 of the JDBC driver.

| OS               | Command                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RHEL             | <p> <b>Important:</b> Using the yum install command to install the MySQL driver package before installing a JDK installs OpenJDK, and then uses the Linux alternatives command to set the system JDK to be OpenJDK. If you intend to use an Oracle JDK, make sure that it is <b>installed</b> before installing the MySQL driver using yum install. If you want to use OpenJDK, you can install the driver using yum.</p> <p>Alternatively, use the following procedure to manually install the driver.</p> <ol style="list-style-type: none"> <li>1. Download the MySQL JDBC driver from <a href="http://www.mysql.com/downloads/connector/j/5.1.html">http://www.mysql.com/downloads/connector/j/5.1.html</a> (in .tar.gz format). As of the time of writing, you can download version 5.1.46 using wget as follows: <pre>wget https://dev.mysql.com/get/Downloads/Connector-J/mysql-connector-java-5.1.46.tar.gz</pre> </li> <li>2. Extract the JDBC driver JAR file from the downloaded file. For example: <pre>tar zxvf mysql-connector-java-5.1.46.tar.gz</pre> </li> <li>3. Copy the JDBC driver, renamed, to /usr/share/java/. If the target directory does not yet exist, create it. For example: <pre>sudo mkdir -p /usr/share/java/ cd my sql-connector-java-5.1.46 sudo cp mysql-connector-java-5.1.46-bin.jar /usr/share/java/ mysql-connector-java.jar</pre> </li> </ol> |
| SLES             | <pre>sudo zypper install mysql-connector-java</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Ubuntu or Debian | <pre>sudo apt-get install libmysql-java</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

### Installing the MySQL JDBC Driver

Install the JDBC driver on the Cloudera Manager Server host, as well as any other hosts running services that require database access.




**Note:** If you already have the JDBC driver installed on the hosts that need it, you can skip this section. However, MySQL 5.6 requires a 5.1 driver version 5.1.26 or higher.

Cloudera recommends that you consolidate all roles that require databases on a limited number of hosts, and install the driver on those hosts. Locating all such roles on the same hosts is recommended but not required. Make sure to install the JDBC driver on each host running roles that access the database.



**Note:** Cloudera recommends using only version 5.1 of the JDBC driver.

| OS     | Command                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RHEL   | <p> <b>Important:</b> Using the yum install command to install the MySQL driver package before installing a JDK installs OpenJDK, and then uses the Linux alternatives command to set the system JDK to be OpenJDK. If you intend to use an Oracle JDK, make sure that it is <a href="#">installed</a> before installing the MySQL driver using yum install. If you want to use OpenJDK, you can install the driver using yum.</p> <p>Alternatively, use the following procedure to manually install the driver.</p> <ol style="list-style-type: none"> <li>1. Download the MySQL JDBC driver from <a href="http://www.mysql.com/downloads/connector/j/5.1.html">http://www.mysql.com/downloads/connector/j/5.1.html</a> (in .tar.gz format). As of the time of writing, you can download version 5.1.46 using wget as follows: <pre>wget https://dev.mysql.com/get/Downloads/Connector-J/mysql-connector-java-5.1.46.tar.gz</pre> </li> <li>2. Extract the JDBC driver JAR file from the downloaded file. For example: <pre>tar zxvf mysql-connector-java-5.1.46.tar.gz</pre> </li> <li>3. Copy the JDBC driver, renamed, to /usr/share/java/. If the target directory does not yet exist, create it. For example: <pre>sudo mkdir -p /usr/share/java/ cd my sql-connector-java-5.1.46 sudo cp mysql-connector-java-5.1.46-bin.jar /usr/share/java/ mysql-connector-java.jar</pre> </li> </ol> |
| SLES   | <pre>sudo zypper install mysql-connector-java</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Ubuntu | <pre>sudo apt-get install libmysql-java</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

### Installing the Oracle JDBC Connector

You must install the JDBC connector on the Cloudera Manager Server host and any other hosts that use a database.

1. Download the Oracle JDBC Driver from the Oracle website. For example, the version 6 JAR file is named ojdbc6.jar.

For more information about supported Java versions, see [Java Requirements](#).

To download the JDBC driver, visit the [Oracle JDBC and UCP Downloads](#) page, and click on the link for your Oracle Database version. Download the ojdbc6.jar file (or ojdbc8.jar, for Oracle Database 12.2).

2. Copy the Oracle JDBC JAR file to /usr/share/java/oracle-connector-java.jar. The Cloudera Manager databases and the Hive Metastore database use this shared file. For example:

```
sudo mkdir -p /usr/share/java
sudo cp /tmp/ojdbc8-12.2.0.1.jar /usr/share/java/oracle-connector-java.jar
sudo chmod 644 /usr/share/java/oracle-connector-java.jar
```

### Proxy Cloudera Manager through Apache Knox

In order to have Cloudera Manager proxied through Knox, there are some steps you must complete.

### Procedure

1. Set the value for `frontend_url`: Cloudera Manager Administration Settings Cloudera Manager Frontend URL :
  - Non-HA value: `https://$Knox_host:$knox_port`
  - HA value: `https://$Knox_loadbalancer_host:$Knox_loadbalancer_port`
2. Set allowed groups, hosts, and users for Knox Proxy: Cloudera Manager Administration Settings External Authentication :
  - Allowed Groups for Knox Proxy: \*
  - Allowed Hosts for Knox Proxy: \*
  - Allowed Users for Knox Proxy: \*
3. Enable Kerberos/SPNEGO authentication for the Admin Console and API: Cloudera Manager Administration Settings External Authentication Enable SPNEGO/Kerberos Authentication for the Admin Console and API : `true`
4. From Cloudera Manager Administration Settings External Authentication , set Knox Proxy Principal: `knox`.

### What to do next

External authentication must be set up correctly. Cloudera Manager must be configured to use LDAP, following the standard procedure for setting up LDAP. This LDAP server should be the same LDAP that populates local users on Knox hosts (if using PAM authentication with Knox), or the same LDAP that Knox is configured to use (if using LDAP authentication with Knox).

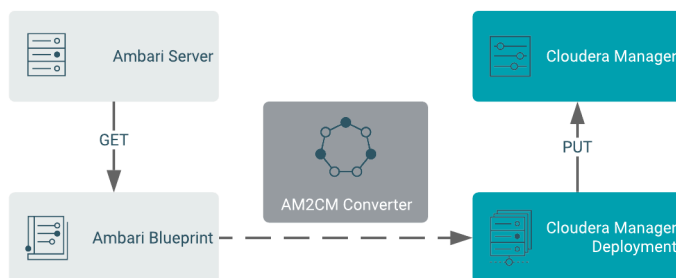
## Transitioning HDP cluster to CDP Private Cloud Base cluster using the AM2CM tool

Use AM2CM tool to transition from Ambari-managed cluster to Cloudera Manager-managed cluster.

The purpose of this tool is to convert the Ambari blueprint to Cloudera Manager Deployment template. Download and export the Ambari blueprint to the AM2CM tool. The AM2CM tool converts the Ambari blueprint to Cloudera Manager Deployment template. Import the converted template to Cloudera Manager, start the services through the Cloudera Manager UI, and validate the cluster.



**Note:** Before you use the AM2CM tool, you must understand the HDP cluster configuration to resolve any edge-case issues that may occur during migration.



## Transitioning HDP 2.6.5 cluster to CDP Private Cloud Base 7.1.x cluster using the AM2CM tool

Use the AM2CM tool to transition from Ambari-managed HDP cluster to Cloudera Manager-managed CDP cluster.

### Transitioning the cluster from HDP to CDP Private Cloud Base

The latest and supported version of the AM2CM tool is AM2CM 2.8.1.0 and also supports the transition from HDP to CDP Private Cloud Base 7.1.8. The AM2CM tool 2.0.2 and higher supports the transition from HDP to CDP Private Cloud Base 7.1.4, 7.1.6, and 7.1.7. For more information, see [Software download matrix](#)



## Related Information

[Configure a Repository for Cloudera Manager](#)

## Transitioning the cluster from HDP to CDP Private Cloud Base using the AM2CM tool

Transition the HDP cluster managed by Ambari to Cloudera Private Cloud Base cluster managed by Cloudera Manager using the AM2CM tool. Here we have taken the AM2CM 2.8.1.0 tool as an example.

### Before you begin

Take a copy of your blueprint before transitioning your cluster. It helps you to understand the changes and differences between HDP 2.6.5 versus HDP 7.1.x versus Cloudera Manager.



#### Note:

If you have more than 200 nodes, you can choose to pre-deploy the CDP parcels. You can pre-deploy the CDP parcels before Cloudera Manager starts managing the CDP Private Cloud Base cluster.

The pre-deployment of parcels can help you in reducing downtime. The downtime is the difference between the shutdown of Ambari-managed HDP Intermediate bits and the start of Cloudera Manager-managed CDP Private Cloud Base.

For more information, see the community article on [How to pre-distribute CDP 7 parcels before migration from HDP 2.6.5](#).

### About this task

Perform the following steps:

### Procedure

1. Log in to the Ambari server.
2. In the new tab, enter the URL: `http://<ambari_ip>:<port>/api/v1/clusters/<cluster-name>?format=blueprint_with_hosts`



**Note:** If there are browser issues, use this API call on CLI:

```
curl -u username:password http://<ambari_ip>:<port>/api/v1/clusters/
<cluster-name>?format=blueprint_with_hosts
```

3. Download the Ambari blueprint and save it in JSON format.
4. Download the AM2CM tool from [Software download matrix](#).
5. Extract the `am2cm-tool-2.8.1.0-1.tar` file to `am2cm-tool-2.8.1.0-1` folder.
6. Set the `JAVA_HOME` variable. For example, `$ export JAVA_HOME= [path to your installed JDK]`
7. The `kafka-broker-ids.ini` file generated must be manually copied to `$am2cm-tool-2.8.1.0-1/conf/` path before running the tool. For more information, see [Extract Kafka Broker ID](#).
8. Navigate to the `am2cm-tool-2.8.1.0-1/conf/user-settings.ini` file and update Parcels, Cluster name, passwords, and JDBC URL information.



#### Note:

- When you are configuring the password fields in the `user-setting.ini` file, you must not enclose the password with double-quotes. You must provide correct passwords. If you use double quotes, then this causes an issue when you import the template into Cloudera Manager.
- For more information on resetting the Ranger Web UI admin's password, see the [Knowledge Base article](#).

```
# Cluster details cluster.name=<Cluster-Name> cluster.displayname=<Cluster-Name>
                                ambari.cluster.name=<Ambari-cluster-Name cm
.rolegroups.enable = false
```

```

# Hive JDBC settings
SERVICE_HIVE_hive_metastore_database_passwo
rd=<DB-Password> SERVICE_HIVE_hive_jdbc_url_override=<JDBC_URL>

# Oozie JDBC settings
SERVICE_OOZIE_oozie_database_password=<DB-Pas
sword> SERVICE_OOZIE_oozie_service_JPAService_jdbc_url=<JDBC_URL>

# Ranger JDBC settings
SERVICE_RANGER_ranger_database_password=<Range
r-DB-Password>
SERVICE_RANGER_rangeradmin_user_password=<Ran
geradmin-user-Password>
SERVICE_RANGER_rangerusersync_user_password=
<Rangerusersync-user-Password>
SERVICE_RANGER_rangertagsync_user_password=<R
angertagsync-user-Password>
SERVICE_RANGER_rangerkeyadmin_user_password=
<Rangerkeyadmin-user-Password>
SERVICE_RANGER_KMS_ranger_kms_master_key_pass
word=<Rangerksmmaster-key-Password>
SERVICE_RANGER_KMS_ranger_kms_database_passw
ord=<Rangerkms-database-Password>

#Knox Settings
SERVICE_KNOX_gateway_master_secret=admin

```

9. Generate the Cloudera Manager Deployment template.



**Note:** If you want skip any pre-upgrade validations of your blueprint, run the following command: `-spu` or `--skip-pre-upgrade`

```

# cd am2cm-tool-2.8.1.0-1
# chmod +x ./am2cm.sh
# ./am2cm.sh -bp /PATH/TO/Amb_blueprint.json -dt /PATH/TO/cm_deployment
_template.json -sv <hdp_version> -tv <cdp_version>
Where -sv <hdp_version> -tv <cdp_version> can be selected from the below
list:
-sv hdp2 -tv 7.1.4
-sv hdp2 -tv 7.1.6
-sv hdp2 -tv 7.1.7
-sv hdp2 -tv 7.1.8

```



**Note:** This generates a `cm_deployment_template.json` file in the `am2cm-tool-2.8.1.0-1` folder.

10. Check for errors in the console or in the `am2cm-tool-2.8.1.0-1/cm_migration.log` logs.

11. Stop HDP services from Ambari.

12. In Ambari UI, disable Auto Start Settings. For more information, see [Disable service auto start settings from Ambari Web](#)

**13.** Import the template into Cloudera Manager using the Cloudera Manager API through browser or CLI.

- Using browser: If you are using browser, you must authenticate to the Swagger UI using the Cloudera Manager full admin credentials.
  - a. Copy the URL in the browser and enter. `http://<CM_HOST>:7180/static/apidocs/ui/index.html#!/ClouderaManagerResource/updateDeployment2`
  - b. Navigate to Cloudera ManagerResource
  - c. Copy the template in body.
  - d. Click Tryit out!
- Using CLI:

```
curl --user admin:admin -k -X PUT -H "Content-Type: application/json" -d
  @cm_deployment_template.json 'http://<CM_HOST>:7180/api/v41/cm/deployment?deleteCurrentDeployment=false'
```

- 14.** In Cloudera Manager Parcel screen, download Cloudera Runtime and distribute the parcel. Cloudera Manager deploys the parcel to the cluster hosts. The default is 10 concurrent hosts. You can adjust this to a maximum of 50 by increasing the setting in the Other Parcel Configurations screen, but must be configured before the template is applied.

7.1.6-1.cdh7.1.6.p0.10506313

Downloaded

Distribute



**Note:**

- a. After the parcels are distributed, the Cloudera Agent unpacks them. Ensure that you have sufficient free space on `/opt/cloudera` to hold both the parcel and its unpacked copy.
- b. By default, the AM2CM tool migrates the Ambari configuration groups as host overrides in Cloudera Manager. To migrate as **Role Groups**, you must enable the `cm.rolegroups.enable=true` parameter. This is an experimental feature and you must validate this in your lower environments before using this in production environment. The migrated Role groups will not have the same configuration group names defined in Ambari. For example, `hdfs-DATANODE-host_group_2`. You can rename this in ClouderaManager after importing the template.

**15. Enable Cloudera Manager TLS (Optional) if you want:**

- Certificate management: Creating certificates, keystores, and truststores.
- Certificate distribution or configuration:
  - Copying keystores & truststores to servers.
  - Configuring services to reference these keystores & truststores.
  - Configuring related TLS properties for service.
- Ensure that you compare Manual TLS and Auto TLS and then proceed. For more information, see [Comparing manual TLS and Auto-TLS](#)
- If you select the Manual TLS option, you must manually configure TLS. For more information, see [Manually configure TLS Encryption for Cloudera Manager](#).
- If you select the Auto TLS option 1, Cloudera Manager handles it independent of any company certificate authority. This is basically creating a private certificate authority that only Cloudera Manager knows about.
  - If you want TLS but you do not have any of the external certificate management infrastructures then you will probably want this.
  - The benefit is you get full automation for the cluster side (management & certificate distribution and configuration) but requires client configuration to trust the private certificate authority. For more information, see [Auto TLS 1](#)
- If you select Auto TLS option 2a, Cloudera Manager handles certificate management based on a company certificate authority. Cloudera Manager generates certificates on your behalf using the certificate authority and performs distribution and configuration for you.
  - If you want TLS and are willing to extend trust from an external certificate authority to Cloudera Manager and allow Cloudera Manager to generate certificates will want this.
  - The benefit is you get full automation for the cluster side (management & certificate distribution and configuration) but requires extending trust to Cloudera Manager. Clients need not require any additional configuration because they would already trust the global company certificate authority. For more information, see [Auto TLS 2](#)
- If you select AutoTLS option 2b, you are only doing certificate distribution and configuration because you are doing certificate management outside of Cloudera Manager and manually loading those certificates into Cloudera Manager's certificate repository.
  - If you want TLS but are unwilling to extend trust from an external certificate authority to Cloudera Manager will want this.
  - The benefit is you get partially automated for the cluster side (certificate distribution and configuration only). Per-host or per-service certificate management done outside of Cloudera Manager and certificates manually uploaded into Cloudera Manager by an admin. For more information, see [Auto TLS 2](#)

**16. After the parcels are deployed on Cloudera Manager, activate the Cloudera Runtime 7.1.x parcels.**

| Cluster 1        |                               |                                                                           |
|------------------|-------------------------------|---------------------------------------------------------------------------|
| Parcel Name      | Version                       | Status                                                                    |
| Cloudera Runtime | 7.1.7-1.cdh7.1.7.p10.17435777 | Distributed <span style="float: right;"><a href="#">Activate</a> ▼</span> |

**17. LZ0 package configuration: If HDP is configured with LZ0 packages then follow the steps given below.**

- a) Go to the Parcels/Parcel Repository & Network Settings page on Cloudera Manager user interface and add Remote Parcel Repository URLs. [URL](#)
- b) In the Parcels screen - Download, Distribute, and Activate “GPLEXTRAS”. If the HDP intermediate bits or Ambari 7.1.x.x has used LZ0 packages, then enable or add the packages to Cloudera Manager. For more information, see the [Configuring Data Compression](#) documentation.

## Post transition steps

You must complete the post transition steps to start the services in Cloudera Private Cloud Base.

The AM2CM tool transitions the component configurations. However, you must configure and perform additional steps to start the services in Cloudera Private Cloud Base.



### Note:

- Review configuration warnings for all the services.
- Review JVM parameters and configurations for all services as some of the JVM parameters and configurations are not transitioned. Most of the heap configurations are transitioned.
- Review the Log4j configurations. Log4j configurations such as logs dir, size, and backup index are transitioned to Cloudera Manager as it uses the default log4j settings of Cloudera Manager.



**Note:** If the cluster is kerberized, you can generate keytabs in Cloudera Manager under Security section - >Kerberos Credentials -> Generate Missing Credentials

## Enable Auto Start setting

Ensure that you enable Auto Start Settings in Cloudera Manager for all the components.

### About this task

If you enable Auto Start Settings in Ambari, the Ambari services are configured to start automatically on system boot. However, post migration to Cloudera Manager, this Auto Start Settings field is enabled only for Zookeeper-Server. For other components, enable this setting manually in Cloudera Manager.

### Procedure

1. In the Search box, search for Automatically Restart Process. Results are displayed for all the components.
2. Select the component from the results displayed.
3. Click the Automatically Restart Process checkbox.

## Kerberos Principal for Cloudera Manager Server

How to create a Kerberos principal for Cloudera Manager Server.

For information on creating Kerberos principal for Cloudera Manager Server, see [Create the Kerberos Principal for Cloudera Manager Server](#).

## ZooKeeper

You must complete the following steps to start the ZooKeeper service.

### Procedure

1. Remove or move the myid file from the ZooKeeper server hosts. The path is `${dataDir}/myid`. For example,
 

```
# mv /hadoop/zookeeper/myid /hadoop/zookeeper/myid.bak
```
2. Start ZooKeeper service from Cloudera Manager

## Delete ZNODES

Remove the ZNode for HDFS, YARN, Hive, HBase, and Oozie

### About this task

If the cluster is configured for Namenode HA, the Failover controllers cannot determine the Active Namenode after the migration. This is because the naming conventions used to identify Namenodes for the HA service are different

for Ambari and Cloudera Manager. You must format the ZooKeeper ZNode for the Failover controller to continue further.

If the cluster is configured for Yarn HA, the Resource Managers cannot elect an Active Resource Manager after the migration. You must clear the previous Znode from ZooKeeper so that YARN replaces it.

If the cluster is configured for Ranger KMS, you must delete the Ranger KMS Znode with superuser privileges. When you restart Ranger KMS, the znode gets created with correct privileges.

You must format the ZooKeeper Znode for the failover controller to continue further.

### Procedure

1. Start ZooKeeper client CLI session from a master node. (Assuming that the skipACL is set to Yes to avoid authentication issues). `zookeeper-client -server <zk_hostname>`
2. Remove the HDFS HA Failover controller Znode  
`deleteall /hadoop-ha`
3. Remove the YARN ZNode  
`deleteall /yarn-leader-election`
4. Remove the Hive ZNode  
`deleteall /hive`  
`deleteall /hiveserver2`
5. Remove the HBase ZNode  
`deleteall /hbase-secure`
6. Remove the Oozie ZNode  
`deleteall /oozie`
7. Remove the Ranger KMS ZNode  
`deleteall /zktsm`

### Ranger

You must change the port number for Ranger port numbers on Cloudera Manager.

After transitioning the cluster, the `ranger.service.https.port` and `ranger.service.http.port` configurations must have different port numbers in Cloudera Manager. Otherwise, Cloudera Manager displays a warning.

1. Start the Ranger server
2. Click Actions
3. Click Setup Ranger Plugin Service

### Ranger Service connection with Oracle database

If you have configured Ranger with Oracle Database as its backend, you must add the property `conf/ranger-admin-site.xml` with the key and value provided below for the Ranger Admin Advanced Configuration Snippet (Safety Valve) in Cloudera Manager: Key: `ranger.jpa.jdbc.preferredtestquery` value: `select 1 from dual;`

### Ranger KMS

You must add the key and value to Ranger KMS property. This procedure is not required if you are on Cloudera Manager 7.4.4 and above and CDP 7.1.5 and above.

You must set `RANGER_KMS_SERVER_role_env_safety_valve` setting to `JAVA_OPTS` in Cloudera Manager. This is needed only for non default zookeeper principal.

1. Log in to Cloudera Manager UI
2. Navigate to Clusters
3. Select the Ranger KMS service

4. Go to Configurations
5. Search for RANGER\_KMS\_SERVER\_role\_env\_safety\_valve and add key JAVA\_OPTS and value -Dzookeeper.sasl.client.username=<zookeeper custom principal name set on the cluster> For example, if the principal is set to custom principal value zk-test then, -Dzookeeper.sasl.client.username=zk-test
6. Start Ranger KMS service

## Add Ranger policies for components on the CDP Cluster

Enable the default ranger policies on the CDP Private Cloud Base cluster.

### Before you begin

The default Ranger policies created on the CDP cluster under cm\_<servicename> are effective after the policies are moved under <clustername>\_<servicename>. For example, in the case of Knox, the default ranger policy for Knox in CDP is CDP Proxy UI and API that is created under the service name cm\_knox which can be moved under cl1\_knox. The <clustername> is the name registered under ambari.

### About this task

To migrate policies from cm\_<servicename> to <clustername>\_<servicename>, perform the following

### Procedure

1. Log in to Cloudera Manager UI
2. Navigate to Clusters
3. Select the Ranger service
4. Click Web UI. This redirects to the Ranger service page.
5. On the Ranger Admin UI, click Service Manager
6. Add the policies available in the table

**Table 3:**

| Service Type | Policy Name           | Existing Groups | Existing Users                  | Existing Permissions | Add new default users to the existing policies | Add new default groups | Add new Permissions |
|--------------|-----------------------|-----------------|---------------------------------|----------------------|------------------------------------------------|------------------------|---------------------|
| Kafka        |                       |                 |                                 |                      |                                                |                        |                     |
| *            | all-topic             |                 | kafka                           |                      | cruisecontrol, streamsmgmgr, streamsrepmgr     |                        | N.A                 |
|              |                       |                 | rangerlookup                    |                      |                                                |                        | N.A                 |
| *            | all - cluster         |                 | kafka                           |                      | cruisecontrol, streamsmgmgr, streamsrepmgr     |                        | N.A                 |
|              |                       |                 | rangerlookup                    |                      |                                                |                        | N.A                 |
| *            | all - transactionalid |                 | kafka                           |                      | cruisecontrol, streamsmgmgr, streamsrepmgr     |                        | N.A                 |
|              |                       |                 | rangerlookup                    |                      |                                                |                        | N.A                 |
| *            | all - delegationtoken |                 | kafka                           |                      | cruisecontrol, streamsmgmgr, streamsrepmgr     |                        | N.A                 |
|              |                       |                 | rangerlookup                    |                      |                                                |                        | N.A                 |
| *            | ATLAS_HOOK            |                 | hive, hbase, storm, spark_atlas |                      | impala, mlgov, nifi                            |                        | N.A                 |

| Service Type | Policy Name                                                                                                                                                     | Existing Groups | Existing Users                   | Existing Permissions | Add new default users to the existing policies                                  | Add new default groups | Add new Permissions                       |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|----------------------------------|----------------------|---------------------------------------------------------------------------------|------------------------|-------------------------------------------|
|              |                                                                                                                                                                 |                 | atlas                            | consume              |                                                                                 |                        | create<br>configure<br>consume<br>publish |
| *            | ATLAS_ENTITIES                                                                                                                                                  |                 | atlas                            | publish              |                                                                                 |                        | create<br>configure<br>publish            |
|              |                                                                                                                                                                 |                 | rangertagsync                    |                      |                                                                                 |                        |                                           |
| Atlas        |                                                                                                                                                                 |                 |                                  |                      |                                                                                 |                        |                                           |
| *            | all - entity-type, entity-classification, entity                                                                                                                |                 | atlas,<br>rangerlookup,<br>admin |                      | beacon,<br>dpprofiler,admin,<br>nifi, admin<br>(remove atlas user from policy)  |                        |                                           |
|              |                                                                                                                                                                 |                 | rangertagsync                    |                      |                                                                                 | public                 |                                           |
|              |                                                                                                                                                                 |                 | atlas                            |                      | (remove atlas user from policy)                                                 |                        |                                           |
|              |                                                                                                                                                                 |                 |                                  |                      | rangerlookup                                                                    |                        | entity-read                               |
| *            | all - relationship-type, end-one-entity-type, end-one-entity-classification, end-one-entity, end-two-entity-type, end-two-entity-classification, end-two-entity | public          | atlas<br>rangerlookup<br>admin   |                      | beacon,<br>dpprofiler,<br>admin, nifi,<br>admin (remove atlas user from policy) |                        |                                           |
|              |                                                                                                                                                                 |                 | atlas                            |                      | (remove atlas user from policy)                                                 |                        |                                           |
| *            | all - atlas-service                                                                                                                                             |                 | atlas,<br>rangerlookup,<br>admin |                      | beacon,<br>dpprofiler,<br>admin, nifi,<br>admin (remove atlas user from policy) |                        | admin-purge,<br>admin-audits              |
|              |                                                                                                                                                                 |                 | atlas                            |                      | (remove atlas user from policy)                                                 |                        |                                           |
| *            | all - type-category, type                                                                                                                                       |                 | atlas,<br>rangerlookup,<br>admin |                      | beacon,<br>dpprofiler,<br>admin, nifi,<br>admin (remove atlas user from policy) |                        |                                           |
|              |                                                                                                                                                                 |                 | atlas                            |                      | (remove atlas user from policy)                                                 |                        |                                           |
|              |                                                                                                                                                                 | public          |                                  |                      |                                                                                 |                        |                                           |



| Service Type | Policy Name                                                                | Existing Groups | Existing Users     | Existing Permissions | Add new default users to the existing policies                                           | Add new default groups | Add new Permissions                                      |
|--------------|----------------------------------------------------------------------------|-----------------|--------------------|----------------------|------------------------------------------------------------------------------------------|------------------------|----------------------------------------------------------|
| *            | all - entity-type, entity-classification, entity, entity-label             |                 | atlas              |                      | beacon, dpprofiler, admin, nifi, admin (remove atlas user from policy)                   |                        |                                                          |
|              |                                                                            | public          | rangertagsync      |                      |                                                                                          |                        |                                                          |
|              |                                                                            |                 |                    | rangerlookup         |                                                                                          |                        | entity-read                                              |
| *            | all - entity-type, entity-classification, entity, entity-business-metadata |                 | atlas              |                      | beacon, dpprofiler, admin, nifi, admin (remove atlas user from policy)                   |                        |                                                          |
|              |                                                                            | public          | rangertagsync      |                      |                                                                                          |                        |                                                          |
|              |                                                                            |                 |                    | rangerlookup         |                                                                                          |                        | entity-read                                              |
| *            | Allow users to manage favorite searches                                    |                 | {USER}             |                      |                                                                                          |                        | entity-read, entity-create, entity-update, entity-delete |
|              | Entity Type = __AtlasUserProfile<br>__AtlasUserSavedSearch                 |                 |                    |                      |                                                                                          |                        |                                                          |
|              | Entity Classification = *                                                  |                 |                    |                      |                                                                                          |                        |                                                          |
|              | Entity ID = {USER} {USER}.*                                                |                 |                    |                      |                                                                                          |                        |                                                          |
| Hive         |                                                                            |                 |                    |                      |                                                                                          |                        |                                                          |
| *            | all - hiveservice                                                          |                 | hive, rangerlookup |                      | hive, beacon, dpprofiler, hue, admin, impala (remove rangerlookup user from policy item) |                        |                                                          |
|              |                                                                            |                 |                    |                      | rangerlookup                                                                             |                        | read                                                     |
| *            | all - global                                                               |                 | hive, rangerlookup |                      | hive, beacon, dpprofiler, hue, admin, impala (remove rangerlookup user from policy item) |                        |                                                          |
|              |                                                                            |                 |                    |                      | rangerlookup                                                                             |                        | read                                                     |

| Service Type | Policy Name                                | Existing Groups    | Existing Users       | Existing Permissions | Add new default users to the existing policies                                              | Add new default groups | Add new Permissions |
|--------------|--------------------------------------------|--------------------|----------------------|----------------------|---------------------------------------------------------------------------------------------|------------------------|---------------------|
| *            | all - url                                  |                    | hive, rangerlookup   |                      | hive, beacon, dpprofiler, hue, admin, impala<br>(remove rangerlookup user from policy item) |                        |                     |
|              |                                            |                    |                      |                      | rangerlookup                                                                                |                        | read                |
| *            | all - database, table, column              |                    | hive, rangerlookup   |                      | hive, beacon, dpprofiler, hue, admin, impala<br>(remove rangerlookup user from policy item) |                        |                     |
|              |                                            |                    |                      |                      | rangerlookup                                                                                |                        | read                |
|              |                                            |                    |                      |                      | {OWNER}                                                                                     |                        | all                 |
| *            | all - database, table, column              |                    | hive, rangerlookup   |                      | hive, beacon, dpprofiler, hue, admin, impala<br>(remove rangerlookup user from policy item) |                        |                     |
|              |                                            |                    |                      |                      | rangerlookup                                                                                |                        | read                |
|              |                                            |                    |                      |                      | {OWNER}                                                                                     |                        | all                 |
| *            | default database tables columns            |                    |                      |                      |                                                                                             | public                 | create              |
|              | Hive Database                              | default            |                      |                      |                                                                                             |                        |                     |
|              | Hive Table                                 | *                  |                      |                      |                                                                                             |                        |                     |
|              | Hive Column                                | *                  |                      |                      |                                                                                             |                        |                     |
| *            | Information_schema database tables columns |                    |                      |                      |                                                                                             | public                 | select              |
|              | Hive Database                              | information_schema |                      |                      |                                                                                             |                        |                     |
|              | Hive Table                                 | *                  |                      |                      |                                                                                             |                        |                     |
|              | Hive Column                                | *                  |                      |                      |                                                                                             |                        |                     |
| HDFS         |                                            |                    |                      |                      |                                                                                             |                        |                     |
| *            | all - path                                 |                    | hadoop, rangerlookup |                      | Remove rangerlookup from policy item                                                        |                        |                     |

| Service Type | Policy Name | Existing Groups | Existing Users | Existing Permissions | Add new default users to the existing policies | Add new default groups | Add new Permissions |
|--------------|-------------|-----------------|----------------|----------------------|------------------------------------------------|------------------------|---------------------|
|              |             |                 |                |                      |                                                | rangerlookup           | read                |

| Knox | Policy Name          | Knox topolgy             | Knox service | Allow conditions - Group | Allow conditions - Permission |
|------|----------------------|--------------------------|--------------|--------------------------|-------------------------------|
|      | CDP Proxy UI and API | cdp-proxy, cdp-proxy-api | *            | public                   | Allow                         |

## Ranger Installation in High Availability with Load Balancer

For clusters that have multiple users and production availability requirements, you may want to configure Ranger high availability (HA) with a load-balancing proxy server to relay requests to and from Ranger.

### About this task

If Ranger Admin High Availability is not configured in Ranger, then you must configure it. For more information, see [Configure Ranger Admin High Availability with a Load Balancer](#).

### Procedure

1. Update the Ranger Load Balancer Address property (`ranger.externalurl`) with the load balancer host URL and port



**Note:** While updating this property, avoid adding an extra trailing slash at the end of the load balancer URL.

2. Click Save Changes
3. In kerberos cluster, create a composite keytab file having HTTP principals of both Ranger Admin instances and Load balancer. For more information, see [Create composite keytab for Ranger HA](#)
4. Update the configurations using Safety Valve for `ranger-admin-site.xml`
  - `ranger.spnego.kerberos.keytab=/etc/security/keytabs/ranger.ha.keytab`
  - `ranger.spnego.kerberos.principal=*`
5. Click Save
6. Restart the Ranger service

### Create composite keytab for Ranger HA

If Kerberos is configured on your cluster, you must create a composite keytab.

### Procedure

1. Log in to the Load balancer node
2. Run `kadmin.local` if KDC server and LB node are same. Otherwise, login to the KDC server hosts and run the `kadmin.local` command
3. Run `mkdir -p /etc/security/keytabs`



**Note:** Run the command `mkdir -p /etc/security/keytabs` on all RANGER\_ADMIN instances node and LB node.

4. Add spnego principal entry of the LB node using `ktadd -norandkey -kt /etc/security/keytabs/ranger.ha.keytab HTTP/hostname@EXAMPLE.COM`
5. Add spnego principal entry of the node where RANGER\_ADMIN is first installed using `ktadd -norandkey -kt /etc/security/keytabs/ranger.ha.keytab HTTP/hostname@EXAMPLE.COM`

6. Add spnego principal entry of the node where RANGER\_ADMIN is next installed using `ktadd -norandkey -kt /etc/security/keytabs/ranger.ha.keytab HTTP/hostname@EXAMPLE.COM`
7. Verify `/etc/security/keytabs/ranger.ha.keytab` having an entry of all the required spnego principals using `klist -kt /etc/security/keytabs/ranger.ha.keytab`
8. `scp /etc/security/keytabs/ranger.ha.keytab` file to other nodes where RANGER\_ADMIN is installed
9. Update permission to `chmod 444 /etc/security/keytabs/ranger.ha.keytab`
10. Update ownership to `chown ranger:hadoop /etc/security/keytabs/ranger.ha.keytab`



**Note:** You must perform these steps each time the Spnego principal is regenerated.

## Set maximum retention days for Ranger audits

You can now update the solr document expiry `ranger.audit.solr.config.ttl` and `ranger.audit.solr.config.delete.trigger` parameters from Ranger configurations in Cloudera Manager and refresh the configurations to get the Solr collection for Ranger audits updated with ttl and delete trigger.

In HDP 2.6.5, the `ranger_audit_max_retention_days` attribute in Ranger Configurations is used to specify the maximum number of days to retain the Ranger Audit records in Ambari Infra solr. By default it is set to 90 days.

1. Log in to Cloudera Manager
2. Navigate to Clusters
3. Select the Ranger service
4. Go to Configurations
5. Search for the `ranger.audit.solr.config.ttl` and `ranger.audit.solr.config.delete.trigger` parameters and update them.

Cloudera Manager displays a refresh configuration notification for the Ranger service.

6. You must refresh the **Configurations** action which updates the parameters in the Ranger audits collection in Solr. For more information, see the [Manual process](#) available.

## HDFS

Perform the following post migration steps.

Configure the ports, check the HA NameNode, custom topology, and review other configurations to ensure that HDFS service is set up.

### Ports

Ensure that you review service configuration warnings and address inconsistencies.

Start all role instances in HDFS service.

### TLS/SSL

If the TLS/SSL is not enabled in the CDP Private Cloud Base cluster, then you must reset the `dfs.data.transfer.protection` configuration.

### About this task

Perform the following steps

### Procedure

1. Log in to Cloudera Manager
2. Navigate to Clusters
3. Select the HDFS service
4. Search for `dfs.data.transfer.protection` and click the click to revert to default option
5. Click Save changes

## HDFS HA

You must check HDFS HA and set Failover controller.

### About this task

Perform the task.

### Procedure

1. Cloudera Manager does not support the comma separated entries. Review `dfs.data.transfer.protection` and `hadoop.rpc.protection` parameters. If SSL is not enabled, then click Undo. This will clear the selections.



**Note:** Privacy option requires SSL.

The screenshot shows the configuration page for 'DataNode Data Transfer Protection'. The parameter is 'dfs.data.transfer.protection' and is set to 'hdfs (Service-Wide)'. There are three radio button options: 'Authentication', 'Integrity', and 'Privacy'. The 'Privacy' option is selected. An 'Undo' button is located in the top right corner of the configuration area.

2. Start HDFS service.  
Failover Controller fails to start. You must continue with step 3 or 4.
3. Format the ZooKeeper Failover controller ZNode. SSH to the Failover controller host And perform the following using the example:
 

```
# cd /var/run/cloudera-scm-agent/process/<xx>-hdfs-FAILOVERCONTROLLER /opt/cloudera/parcels/CDH/bin/hdfs --config /var/run/cloudera-scm-agent/process/<xx>-hdfs-FAILOVERCONTROLLER/ zkfc -formatZK
```

  - a) Check the output in the log for success. If the error is: ``ERROR tools.DFSZKFailoverController: DFSZKFailOverController exiting due to earlier exception java.io.IOException: Running in secure mode, but config doesn't have a keytab``.
  - b) The ZooKeeper nodes are not created. You must manually create the ZNodes for the Failover controllers to start and allow an Active Namenode to be elected.
  - c) With the zookeeper-client, open a ZooKeeper Shell ``zookeeper-client -server <a_zk_server>``
  - d) Create the required ZNodes `create /hadoop-ha create /hadoop-ha/<namenode_namepace>`.
4. This is an alternative step to the above step 3.
  - a) Log in to Cloudera Manager.
  - b) Navigate to Clusters
  - c) Click HDFS
  - d) Go to Configuration tab
  - e) Under Filters > Scope, click Failover Controller
  - f) Navigate to the Failover Controller instance
  - g) Initialize Automatic Failover Znode.
5. Start HDFS service again.

### Custom Topology

Custom topology configuration in Cloudera Manager.

In the Ambari HDFS service, if `net.topology.script.file.name` is configured with custom topology file, then, in Cloudera Manager > HDFS Configuration page, you must configure `net.topology.script.file.name` with the same file.

### Add Balancer Role to HDFS

Use HDFS Balancer for balancing the data. This section is not applicable if you are upgrading to CDP Private Cloud Base 7.1.7.

The HDFS Balancer is a tool for balancing the data across the storage devices of a HDFS cluster. For more information on HDFS Balancer, see *Balancing data across an HDFS cluster*.

### Related Information

[Balancing data across an HDFS cluster](#)

### Other review configurations for HDFS

1. Review `dfs.client.failover.proxy.provider.mycluster` value. If this parameter has a non-default value, then configure this parameter in Cloudera Manager. This configuration is not migrated from the Ambari managed HDP cluster to Cloudera Manager.
2. Review and add the cloud service related configurations as those are not transitioned to Cloudera Manager.

### Configuring HDFS properties to optimize log collection

CDP uses “out\_webhdfs” Fluentd output plugin to write records into HDFS, in the form of log files, which are then used by different Data Services to generate diagnostic bundles. Over time, these log files can grow in size. To optimize the size of logs that are captured and stored on HDFS, you must update certain HDFS configurations in the `hdfs-site.xml` file using Cloudera Manager.

### Procedure

1. Log in to Cloudera Manager as an Administrator.
2. Go to Clusters HDFS service Configuration .
3. Select the Enable WebHDFS (`dfs_webhdfs_enabled`) option.
4. Add the following lines in the HDFS Service Advanced Configuration Snippet (Safety Valve) for `hdfs-site.xml` field by clicking View as XML to enable append operations:

```
<property>
  <name>dfs.support.append</name>
  <value>true</value>
</property>

<property>
  <name>dfs.support.broken.append</name>
  <value>true</value>
</property>
```

5. Click Save Changes.
6. Restart the HDFS service.
7. Restart your CDP cluster.

### Related Information

[Fluentd documentation](#)

## Solr

You must create a Ranger Plug-in audit directory and HDFS Home directory and start the Solr service.

On the Cloudera Manager UI:

1. Initialize Solr
2. Create Ranger Plug-in audit directory
3. Create HDFS Home directory
4. Start Solr service

### Restore Solr collections on CDP cluster

To restore the data on Solr managed by Cloudera Manager, you must start HDFS and Zookeeper only. This ensures no operations are performed by services which can push data to Solr.

### About this task

To restore the data on Solr managed by Cloudera Manager, perform the following steps:

### Procedure

1. After migrating to CDP Private Cloud Base, you must update the `infra_solr` section ini file by collecting properties from Solr on Cloudera Manager. (keep a backup of the existing file);

- protocol: Solr protocol (http if TLS is not enabled https if TLS is enabled)
- hosts: Solr server hosts
- zk\_connect\_string: Solr Zookeeper connection string, should be available as part of process Environment Variables
- znode: Solr Znode configured, default value is `/solr_user : Solr service`
- user: default value is `solr`
- port: Solr Server port, default 8983 and 8985 when TLS is enabled for Solr
- keytab: Absolute path for Solr keytab
- principal: Solr Service principal
- zk\_principal\_user: Zookeeper principal

2. After updating the ini configuration file, restore the solr collections using [AM2CM 2.0.0.0 tool](#)

- To change ownership of the Solr data in HDFS

```
python restore_collections.py --ini-file $CONFIG_INI_LOCATION --action change-ownership-in-hdfs
```

- To delete newly created default Cloudera Manager Solr collections

```
python restore_collections.py --ini-file $CONFIG_INI_LOCATION --action delete-new-solr-collections
```

- To restore original collections

```
python restore_collections.py --ini-file $CONFIG_INI_LOCATION --action full-restore
```

3. If you do not want to update the ini file you can have the same parameters handy and pass them as arguments to the above three commands.

- ```
python restore_collections.py --ini-file $CONFIG_INI_LOCATION --action change-ownership-in-hdfs --protocol=<Solr service protocol> --hosts=<Solr Server Hosts comma separated> --zk-connect-string=<Zookeeper Connection String> --znode=/solr-infra --user=solr --port=8985 --keytab=<full path to solr.keytab> --principal=<Solr principal> --zk-principal-user=<Zookeeper principal>
```

- ```
python restore_collections.py --ini-file $CONFIG_INI_LOCATION --action delete-new-solr-collections --action change-ownership-in-hdfs --protocol=<Solr service protocol> --hosts=<Solr Server Hosts comma separated> --zk-connect-string=<Zookeeper Connection String> --znode=/solr-infra --user=solr --port=8985 --keytab=<full path to solr.keytab> --principal=<Solr principal> --zk-principal-user=<Zookeeper principal>
```

- ```
python restore_collections.py --ini-file $CONFIG_INI_LOCATION --action full-restore --action change-ownership-in-hdfs --protocol=<Solr service protocol> --hosts=<Solr Server Hosts comma separated> --zk-connect-string
```

```
=<Zookeeper Connection String> --znode=/solr-infra --user=solr --port=89
85 --keytab=<full path to solr.keytab> --principal=<Solr principal> --zk
-principal-user=<Zookeeper principal>
```

## Kafka

You must cleanup metadata on broker hosts after migrating the Ambari-managed HDP cluster to Cloudera Private Cloud Base. If you are upgrading to CDP Private Cloud Base 7.1.7, you can ignore step 3, step 4, and step 5.

### About this task

Cleanup Metadata on Broker Hosts

### Procedure

1. On each Broker host, remove `$log.dirs/meta.properties` file from Kafka broker hosts. For example,
 

```
#mv /grid/0/kafka-logs/meta.properties /tmp.
```

 Remove `/${log.dirs}/meta.properties` file from Kafka broker.  
 If the Kafka `log.dirs` property points to `/kafka-logs`, then the command is
 

```
#mv /kafka-logs/meta.properties /tmp
```
2. Set the Kafka port value. For more information, see [Change Kafka Port Value](#).
3. Enable `kerberos.auth.enable`. For more information, see [Kafka cluster Kerberos](#).
4. In the Ambari kafka configuration, get the `zookeeper.connect` configuration value. Find the path defined at the end after port number and update the path at Cloudera Manager Kafka `zookeeper.chroot` configuration. For example, if the path is `zookeeper.connect=hostname1:port1,hostname2:port2,hostname3:port3/chroot/path`, then the path is `/chroot/path`. Another example, if the path is `zookeeper.connect=hostname1:port1,hostname2:port2,hostname3:port3`, then the path is `/`. Also, the default path is `/`.
5. Update the broker IDs:
  - a. Log in to Cloudera Manager
  - b. Navigate to Clusters
  - c. Select the Kafka service
  - d. Navigate to the Configurations tab
  - e. Search for `broker.id` and update the IDs for each hostname using the output of the third step of [Extract broker ID Procedure 2](#)
6. Start the Kafka service.

### Change Kafka port value

Change the Kafka JMX port value if the Ranger KMS port is conflicting with Kafka JMX port.

The Ranger KMS `ranger.service.https.port` port value is migrated from HDP. In Cloudera Manager, if the Kafka `jmx_port` port value is the same as Ranger KMS port value from HDP, then you must change the Kafka JMX port to a different value in Cloudera Manager.

### Kafka cluster Kerberos

If the CDP cluster is kerberised, then you must set this parameter. This setting is applicable only if you are upgrading from HDP 2.6.5 to CDP Private Cloud Base 7.1.6.

### Procedure

1. Log into Cloudera Manager UI
2. Navigate to Clusters
3. Select the Kafka service
4. Go to Configurations



5. Search for `kerberos.auth.enable` and enable the parameter

### Unsetting Kafka Protocol version

You must unset the Kafka protocol version from Cloudera Manager after the upgrade is complete. This allows you to use the updated Kafka version using Cloudera Manager.

On the HDP cluster, using Ambari, you have set the Kafka protocol versions to the older Kafka version to make sure that all the messages are stored in a manner that Kafka can handle. After upgrading to the CDP cluster, the messages are still stored according to the older version of Kafka. This gives you the ability to roll it back if something goes wrong.

After the Kafka Protocol version properties are unset using Cloudera Manager, Kafka will use its current version (the new and updated Kafka version) and newly stored messages will not be handled by the older Kafka version.



**Note:** Do not unset the older Kafka protocol version until you have completed your upgrade testing. Unsetting the version will prevent rollback to the older Kafka release.

To unset the Kafka Protocol version on Cloudera Manager, perform the following steps:

1. Remove the following properties from the Kafka Broker Advanced Configuration Snippet (Safety Valve) configuration property:
  - `Inter.broker.protocol.version`
  - `log.message.format.version`
2. Save your changes.
3. Perform a rolling restart:
  - Select the Kafka service.
  - Click Actions > Rolling Restart.
  - In the pop-up dialog box, select the options you want and click Rolling Restart.
  - Click Close after the command has finished.

## YARN

Perform the following post migration steps.

### Start job history

You must start job history after migration.

### About this task

To start job history after migration

### Procedure

1. Log in to Cloudera Manager
2. Navigate to Clusters
3. Go to YARN
4. Under the Actions drop-down, click Create Job History Dir
5. Create Job History Dir dialog box appears, click Create Job History Dir.  
Successfully created HDFS directory. message appears

### Yarn Mapreduce framework jars

### About this task

Perform the following steps

### Procedure

1. Go to Cloudera Manager UI.
2. Navigate to YARN
3. Under Actions drop-down, click Install Yarn Mapreduce framework jars

### YARN NodeManager

Set the file permission for Cloudera Manager deployment.

If `yarn.nodemanager.linux-container-executor.group` setting is set to `hadoop` in your HDP cluster, you must reset `yarn.nodemanager.linux-container-executor.group` setting to `yarn` in Cloudera Manager.

If this parameter is not set correctly,

```
Issue Yarn NodeManagers fail to start with error
```

error appears.

1. Log into Cloudera Manager UI.
2. Navigate to Clusters
3. Select the YARN service
4. Go to Configurations
5. Search for `yarn.nodemanager.linux-container-executor.group` and reset this to `yarn`

### YARN NodeManager CGroups

Set the file permission for Cloudera Manager deployment. This setting is applicable only if you are upgrading from HDP 2.6.5 to CDP Private Cloud Base 7.1.6.

### About this task

You must set `yarn.nodemanager.linux-container-executor.cgroups.hierarchy` property to `/hadoop-yarn` in Cloudera Manager.

### Procedure

1. In Cloudera Manager, navigate to `YARN Configurations`.
2. Find the CGroups Hierarchy (`yarn.nodemanager.linux-container-executor.cgroups.hierarchy`) property.
3. Set the property to `/hadoop-yarn`.
4. Click Save Changes.

### Reset ZNode ACLs

If you have a kerberized cluster, you manually have to reset the ZNode ACLs after upgrading from HDP to CDP.

### About this task

ZNode ACLs have to be reset only for kerberized cluster. If you have a non-kerberized cluster, skip this post-upgrade task.



#### Note:

If the cluster is heavily used and the `yarn.resourcemanager.max-completed-applications` is kept high (10k-100k) resetting ACLs on YARN ZooKeeper nodes can take a large amount of time to finish.

### Procedure

1. Log into Cloudera Manager UI.
2. Select the YARN service.
3. Click Actions.
4. Select the Reset ACLs on YARN Zookeeper nodes action.

### Placement rules evaluation engine

If you are upgrading to Cloudera Runtime 7.1.6 or higher you need to learn about the new placement rules evaluation engine.

This list is only an overview of the most prominent changes regarding placement rules evaluation. For a more detailed guide, see [Managing placement rules](#).

### JSON-based format

From Cloudera Runtime 7.1.6 placement rules are described in a JSON-based format. The new engine has full backward compatibility with the old format as well but only if you use safety valves which you must not do. Cloudera recommends managing your placement rules through the YARN Queue Manager UI. The format conversion is automatically done at the first time you use the YARN Queue Manager UI in CDP Private Cloud Base 7.1.6 or higher. For more information about the automatic placement rules format conversion, see [Converting old mapping rule format to JSON-based placement rule format](#).

### Enabling dynamic queues

The new placement rules evaluation engines support the create flag which was also introduced in Cloudera Runtime 7.1.6. Non-existing queues are only created if dynamic queue creation is enabled when the applicable placement rule is created. That means in order to enable dynamic child queues you have to perform two steps:

1. Enable the dynamic auto child creation feature for the queue that you want to use as a parent queue for the dynamic child queues.
2. Ensure that the Create the target queue if it does not exist? property is checked when the applicable placement rule is created.

### Fallback action configuration

With the new placement rules and evaluation engine you can define the action that should happen if the target queue of a placement rule does not exist or it cannot be created. This action is called the fallback action. Fallback action can be configured for each placement rule by the user and it can have the following values:

- Skip: Ignore the current rule and proceed to the next.
- PlaceDefault: Place the application to the default queue root.default (unless it is overridden to something else). This is how Capacity Scheduler worked with the old mapping rules.
- Reject: Reject the submission.

### Placement rule policies

Placement rule policies were introduced as a more user friendly solution for mapping rule creation. Placement rule policies cover the functionality of the placement rules and provide shortcuts for the most common use cases.

All legacy Capacity Scheduler mapping rule options are now available as policy. They provide the same functionality without having to manually write rules. For example, the `u:%user:%primary_group.%user` rule can be achieved using the `primaryGroupUser` policy. However, a custom policy option is available that allows you to define the target queue manually without using any predefined placement rule policy.

### Converting old mapping rule format to JSON-based placement rule format

Reviewing how the old colon-separated mapping rule format is converted into the new JSON-based mapping rules format can help you understand the automatic placement rule conversion.

### Placement Rules

From Cloudera Runtime 7.1.6 placement rules are described in a JSON-based format. This format conversion is automatically done at the first time you use the YARN Queue Manager UI in CDP Private Cloud Base 7.1.6 or higher. If you use Cloudera Runtime 7.1.6 or higher you must use the new JSON-based placement rule format.



**Note:** You cannot manage the new JSON-based placement rules by directly editing them in the `capacity-scheduler.xml` configuration file. Instead, you use the YARN Queue Manager UI to manage placement rules. The table is only provided so that you can better understand how the automatic conversion happens.

**Table 4: Mapping rule conversion into the new JSON-based placement rule format**

| Old mapping rule                                      | New JSON-based placement rule  |
|---|--|
| <code>u:username:root.user.queue</code>               | <pre>{   "type": "user",   "matches": "username",   "policy": "custom",   "customPlacement": "root.user.queue",   "fallbackResult": "placeDefault" }</pre> |
| <code>u:%user:%user</code>                            | <pre>{   "type": "user",   "matches": "*",   "policy": "user",   "fallbackResult": "placeDefault" }</pre>  |
| <code>u:%user:root.parent.%user</code>                | <pre>{   "type": "user",   "matches": "*",   "policy": "user",   "parentQueue": "root.parent",   "fallbackResult": "placeDefault" }</pre>                  |
| <code>u:%user:%primary_group</code>                   | <pre>{   "type": "user",   "matches": "*",   "policy": "primaryGroup",   "fallbackResult": "placeDefault" }</pre>  |
| <code>u:%user:%primary_group.%user</code>             | <pre>{   "type": "user",   "matches": "*",   "policy": "primaryGroupUser",   "fallbackResult": "placeDefault" }</pre>                                      |
| <code>u:%user:root.groups.%primary_group.%user</code> | <pre>{   "type": "user",   "matches": "*",   "policy": "primaryGroupUser",   "parentQueue": "root.groups",   "fallbackResult": "placeDefault" }</pre>      |
| <code>u:%user:%secondary_group</code>                 | <pre>{   "type": "user",   "matches": "*",   "policy": "secondaryGroup",   "fallbackResult": "placeDefault" }</pre>  |
| <code>u:%user:%secondary_group.%user</code>           | <pre>{   "type": "user",</pre>   |

### Application User match

It's worth noting that %application:%application requires a user type matcher. It is because internally, the "\*" is interpreted only for users. If you set the type to application, then the "\*" means to match an application which is named "\*".

### YARN owner permission

You must change the owner of the Yarn user. If you want to add a VPC or Compute cluster after migration, you must make changes for /user/yarn in HDFS.

### Procedure

1. SSH to a cluster node
2. You must have Sudo permission to perform this.
3. Change the owner of /user/yarn to hdfs:supergroup yarn:hadoop with the command: `sudo -u hdfs hdfs dfs -chown hdfs:supergroup /user/yarn`
4. Change the permissions of /user/yarn to drwxr-xr-x with the command: `sudo -u hdfs hdfs dfs -chmod 755 /user/yarn`

### YARN mapreduce paramater

Set the mapreduce.shuffle.ssl.enabled parameter to true. This section is not applicable if you are upgrading to CDP Private Cloud Base 7.1.7.

If the Tez service has the tez.runtime.shuffle.ssl.enable parameter set to true, you must enable the mapreduce.shuffle.ssl.enabled parameter set to true in the YARN service in Cloudera Manager.

1. Log in to Cloudera Manager UI.
2. Navigate to Clusters
3. Select the YARN service
4. Go to Configurations
5. Search for mapreduce.shuffle.ssl.enabled and set to true

## Spark

You must initialize a few directories for the Spark service.

### Before you begin

Run the following directory options:

### Procedure

- Create Spark User Dir
- Create Spark History Log Dir

- Create Spark Driver Log Dir

CDP\_01

The screenshot shows the Cloudera Manager interface for a Spark service. At the top, there is a green checkmark icon, a Spark logo, and the text 'Spark'. Below this, there are tabs for 'Status', 'Instances', and 'Configuration'. The 'Status' tab is selected. A dropdown menu labeled 'Actions' is open, showing various options: Start, Restart, Stop, Add Role Instances, Rename, Delete, Enter Maintenance Mode, Refresh History Server, Create Spark User Dir, Create Spark History Log Dir, Create Spark Driver Log Dir, and Deploy Client Configuration. Below the dropdown, there is a 'Health Tests' section with a 'History Server Health' test showing a green checkmark and the text 'The SPARK\_YARN\_HIS...'. Below that is a 'Status Summary' section with a table listing 'Gateway', 'History Server', and 'Hosts'. At the bottom, there is a 'Health History' section with a 'History Server Health' entry showing a green checkmark and the text '1 Resource Good'.



## Livy2

Review and configure Livy2 parameters.

Livy server heap/memory parameters are not migrated automatically when migrating Ambari-managed HDP cluster to Cloudera Private Cloud Base. You must review and configure these parameters in Cloudera Manager.

## Enabling Spark on YARN for Atlas

Post-migration, Spark hook for Atlas is not enabled by default on AM2CM migrated clusters. you must manually enable Spark hook.

1. Log into Cloudera Manager instance
2. Navigate to Clusters
3. Select the Spark on YARN service
4. Select Configurations tab
5. Check the Enable Spark Lineage checkbox
6. Select Atlas Dependency
7. Click Save Changes

## Enabling SAC manually on Spark

Post-migration to CDP Private Cloud Base, you must manually enable SAC on Spark and Hive.

1. Log in to your Cloudera Manager instance
2. Navigate to Clusters
3. Select the Spark service.
4. Select Configuration tab
5. Search for Atlas Service and select the atlas check box
6. Click Save Changes
7. Navigate to Clusters
8. Select the Hive service
9. Select Configuration tab
10. Search for Atlas Service and select the atlas check box
11. Click Save Changes
12. Restart the affected Services.

## Tez

Install the Tez tar files on HDFS.

### About this task

From the Tez cluster service, install Tez tar files on HDFS used by Hive use and then deploy Client Configuration.

### Procedure

1. Run Upload Tez tar file to HDFS

## 2. Deploy Client Configuration.



**Note:** You must provide sufficient permissions to HDFS or Ranger to install Tez and deploy configuration.

CDP\_01

tez

Actions ▾

- Add Role Instances
- Rename
- Delete
- Enter Maintenance Mode
- Upload Tez tar file to HDFS
- Deploy Client Configuration
- Download Client Configuration

Instances Config Quick Link

Search

Filters

- STATUS
  - None
- COMMISSION
- MAINTENANCE MODE
- RACK ID
- ROLE GROUP
- ROLE TYPE

Selected

atus Re

G

G

G

G

G

## Hive

Perform the post migration steps. If you are expediting the Hive upgrade process and modified the upgrade process to skip materializing every table in the metastore, you need to modify the Hive Strict Metastore Migration (HSMM) process by running the Hive Upgrade Check tool and provided scripts.

### Related Information

[Hive Post-Upgrade Tasks](#)

### Setting up Hive metastore for Atlas

As Administrator, you might plan to recommend Atlas for Hive metadata management and data governance. You have to check that Hive metastore for Atlas is set up, so users can build catalogs of data assets, classify, and govern the assets. If Atlas is not set up you learn how to do so. This section is not applicable if you are upgrading to CDP Private Cloud Base 7.1.7.

### About this task

In this task, you set the name of the Atlas service for Hive metastore to use.

### Procedure

1. In , click Clusters Hive Configurations .
2. Search for Atlas Service.
3. Choose a method based on the results of your search:
  - If finds the Atlas Service, check the checkbox to enable the Hive Metastore hook in your instance.
  - If does not find the Atlas Service, in Hive Service Advanced Configuration Snippet (Safety Valve) for atlas-application properties, enter an XML snippet in the value element that provides the name of your Atlas service, myatlasservice in the example below.

```
<property>
  <name>atlas_service</name>
  <value>myatlasservice</value>
</property>
```

4. Save changes.
5. Restart the Hive metastore service.

### Identifying and fixing invalid Hive schema versions

As Administrator, after upgrading from Ambari-managed HDP to Cloudera Private Cloud Base, you need to identify Hive metastore operations that might fail due to Hive schema version incompatibility.

### About this task

Incompatibility might exist if the upgrade process failed to make schema updates. You need to turn on the Hive Metastore Schema validation process for the metastore during the migration of your workloads to CDP. The Hive metastore captures any schema updates that occur during the upgrade, and displays issues in the Hive metastore logs. With this information, you can use the Apache Hive Schema tool to fix any problems.

### Procedure

1. In Cloudera Manager, click Clusters HIVE Configuration .
2. Check the hive.metastore.server.max.message.size.

Max Message Size for Hive MetaStore

hive.metastore.server.max.message.size

Hive Metastore Server Default Group Undo

100 MiB

3. Set `hive.metastore.server.max.message.size` to the recommended value: 10% of the value of your Java heap size for Hive Metastore Server in bytes, but no more than 21478364. Recommended value: 214748364
4. Click `Clusters HIVE Configuration`, and search for schema.
5. Check `Strict Hive Metastore Schema Validation` to set `hive.metastore.schema.validation` to true.
6. Check the Hive metastore logs and set a compatible metastore schema for the current Hive version using the Apache Hive Schema Tool.

### Fixing statistics

Upgrading or migrating from Hive 1 or Hive 2 to Hive 3 might result in missing statistics. In Hive 3, these missing statistics, when detected by the cost-based optimizer (CBO), could cause datasets to be disregarded. As Data Engineer, you need to fix these statistics after upgrading.

### Procedure

1. Run `DESCRIBE FORMATTED <table>`, and check the value of `numrows`.  
If the value is 0, you must fix statistics.
2. Run `ANALYZE` on the tables and columns to fix the statistics.

```
ANALYZE TABLE credit_card_01.cc_acct COMPUTE STATISTICS[FOR COLUMNS];
```

### Advanced configuration snippet (Safety valve)

You must set the Hive service Advanced Configuration Snippet (Safety Valve) for Ranger. This setting is applicable only if you are upgrading from HDP 2.6.5 to CDP Private Cloud Base 7.1.6.

### Procedure

1. Log in to Cloudera Manager UI
2. Navigate to `Clusters`
3. Select the Hive service
4. Go to `Configurations`
5. Search for `Hive Service Advanced Configuration Snippet (Safety Valve)` for `ranger-hive-security.xml`
6. Add the `ranger.plugin.hive.service.name=<hive-repo-name>` configuration

### Example

```
ranger.plugin.hive.service.name=c11_hive
```

### Remove Hive Ranger property

This section is not applicable if you are upgrading to CDP Private Cloud Base 7.1.7.

From the Hive service, navigate to the `Advanced Configuration Snippet (Safety Valve)` option and you must remove the `org.apache.ranger.authorization.hive.authorizer.RangerHiveAuthorizerFactory` property.

## HBase

Post-migration, Atlas hook for HBase is not enabled by default on AM2CM migrated clusters. you must manually enable Atlas hook.

To enable the Hive Metastore hook in your Cloudera Manager instance:

1. Log in to Cloudera Manager
2. Select the HBase service
3. Select `Configurations` tab
4. Search for `Enable Atlas Hook` that the HBase service depends on and check this.
5. Select `Save Changes`

## HBase RegionServer heap size

Post transition to CDP Private Cloud Base 7.1.7 SP2, you must adjust the HBase RegionServer heap size.

The change introduced in HBASE-26036/HBASE-26777 adds an extra copy of cells into the heap, resulting in out of memory on the region servers if the `hbase_regionserver_java_heapsize` is not set correctly. Also, if the region server is killed because of out of memory issue under a large amount concurrent PUTs with multi-columns and scan the recent result. This issue occurs only during the two stage upgrade from HDP 2.6.5 to CDP Private Cloud Base 7.1.7 SP2.

Workaround: After transitioning to CDP Private Cloud Base 7.1.7 SP2, you must adjust the value of `hbase_regionserver_java_heapsize` to two times more than the original value used in the HDP 2.6.5 cluster.

1. Log in to Cloudera Manager.
2. Select the HBase service.
3. Select Configurations tab.
4. Search for `hbase_regionserver_java_heapsize` and adjust the value.
5. Click Save Changes.

## Installing dependencies for Hue

You must install the `psycopg2` Python package for PostgreSQL-backed Hue and MySQL clients for MariaDB and MySQL databases depending on your operating systems.



**Note:** This task is applicable only if you are upgrading to 7.1.8 and higher. If you are upgrading to 7.1.7 or lower, then you can skip this task.

If you are using Oracle as a backend database for Hue, then review [Using Oracle database with Hue](#) to ensure that Hue connects to your database.

### Installing Python 3.8

Certain services, such as Hue, in CDP 7.1.8 and higher use Python 3.8. You must install Python 3.8 on all the hosts running the affected services after you have installed Cloudera Manager and before adding the services to your cluster.



**Note:** This task is applicable only if you are installing or upgrading to CDP 7.1.8 and higher. If you are installing or upgrading to CDP 7.1.7 or lower, then you can skip this task.



**Attention:** Installing Python 3.8 is mandatory if you want to use Hue.

Ubuntu 20 comes preinstalled with Python 3.8. You must install Python 3.8 manually on CentOS 7, RHEL 8, SLES 12, and Ubuntu 18.

### Installing Python 3.8 on CentOS 7 for Hue

You must install Python 3.8 on all hosts on which you want to run the Hue service after installing Cloudera Manager and before adding the services to your cluster.

### Before you begin

Install the necessary developer tools such as `gcc` and `make` on your system.

Install the following packages before installing Python 3.8:

- `openssl-devel`
- `bzip2-devel`
- `libffi-devel`
- `zlib-devel`

```
yum install gcc openssl-devel bzip2-devel libffi-devel zlib-devel -y
```



**Note:** You cannot install the openssl-devel package on a FIPS cluster.

### Procedure

1. SSH into the host system as a root user.
2. Download Python 3.8 and decompress the package by running the following commands:

```
cd /opt
curl -O https://www.python.org/ftp/python/3.8.12/Python-3.8.12.tgz
```

```
tar -zxvf Python-3.8.12.tgz
```

3. Change directory to where you decompressed the Python 3.8 package:

```
cd /opt/Python-3.8.12
```

4. Install Python 3.8 as follows:

```
./configure --enable-shared --prefix=[***CUSTOM-INSTALL-PATH***]
```



**Attention:** By default, Python could be installed in any one of the following locations:

- /usr/bin
- /usr/local/python38/bin
- /usr/local/bin
- /opt/rh/rh-python38/root/usr/bin

If you are installing Python 3.8 in any other location, then you must specify the path using the `--prefix` option. You must also create a symbolic link pointing to `/usr/bin` or to `/usr/local/bin` by running the following command:

```
ln -s [***SOURCE***] [***DESTINATION***]
```

For example:

```
ln -s [***CUSTOM-PYTON-PATH***]/python3.8 /usr/bin/python3.8
```

The `--enable-shared` option is used to build a shared library instead of a static library.

## 5. Build Python 3.8 as follows:

- a) Run the make command to compile the files:

```
make
```

- b) Run the following command to put the compiled files in the default location or in the custom location that you specified using the `--prefix` option:

```
make install
```

- c) Copy the shared compiled library files (`libpython3.8.so`) to the `/lib64/` directory:

```
cp --no-clobber ./libpython3.8.so* /lib64/
```

The `--no-clobber` option is used to prevent overwriting files.

- d) Change the permissions of the `libpython3.8.so` files as follows:

```
chmod 755 /lib64/libpython3.8.so*
```

If you see an error such as “error while loading shared libraries: `libpython3.8.so.1.0`: cannot open shared object file: No such file or directory”, then run the following command:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/local/lib/
```

6. Change the permission of the Python 3.8 installation directory to `755` so that Hue and its related services can leverage the binary and the site packages, as follows:

```
chmod -R 755 /usr/local/lib/python3.8
```

### What to do next

(For Hue) If you have installed Python 3.8 at a custom location, then you must append the custom path in Cloudera Manager Clusters Hue Configuration Hue Service Environment Advanced Configuration Snippet (Safety Valve) separated by colon (`:`) as follows after and restart the Hue service:

Key: `PATH`

Value: `[**CUSTOM-INSTALL-PATH**]:/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin`

### Installing Python 3.8 on RHEL 8 for Hue

You must install Python 3.8 on all hosts after installing Cloudera Manager and before adding the services to your cluster.

### Before you begin

Install the necessary developer tools such as `gcc` and `make` on your system.

Install the following packages before installing Python 3.8:

- `openssl-devel`
- `bzip2-devel`
- `libffi-devel`
- `zlib-devel`

```
yum install gcc openssl-devel bzip2-devel libffi-devel zlib-devel -y
```



**Note:** You cannot install the `openssl-devel` package on a FIPS cluster.

You must also install the following Python 3.8 packages and libraries:

- python38-devel
- python38-libs
- python38-setuptools
- python38-pip
- python38-pip-wheel



**Note:** Python 3.8 is distributed with Red Hat Enterprise Linux version 8.

### Procedure

1. SSH into the host system as a root user.
2. Download Python 3.8 and decompress the package by running the following commands:

```
cd /opt
curl -O https://www.python.org/ftp/python/3.8.12/Python-3.8.12.tgz

tar -zxvf Python-3.8.12.tgz
```

3. Change directory to where you decompressed the Python 3.8 package:

```
cd /opt/Python-3.8.12
```

4. Install Python 3.8 as follows:

```
./configure --enable-shared --prefix=[**CUSTOM-INSTALL-PATH**]
```



**Attention:** By default, Python could be installed in any one of the following locations:

- /usr/bin
- /usr/local/python38/bin
- /usr/local/bin
- /opt/rh/rh-python38/root/usr/bin

If you are installing Python 3.8 in any other location, then you must specify the path using the `--prefix` option. You must also create a symbolic link pointing to `/usr/bin` or to `/usr/local/bin` by running the following command:

```
ln -s [**SOURCE**] [**DESTINATION**]
```

For example:

```
ln -s [**CUSTOM-PYTON-PATH**]/python3.8 /usr/bin/python3.8
```

The `--enable-shared` option is used to build a shared library instead of a static library.



## 5. Build Python 3.8 as follows:

- a) Run the make command to compile the files:

```
make
```

- b) Run the following command to put the compiled files in the default location or in the custom location that you specified using the `--prefix` option:

```
make install
```

- c) Copy the shared compiled library files (`libpython3.8.so`) to the `/lib64/` directory:

```
cp --no-clobber ./libpython3.8.so* /lib64/
```

The `--no-clobber` option is used to prevent overwriting files.

- d) Change the permissions of the `libpython3.8.so` files as follows:

```
chmod 755 /lib64/libpython3.8.so*
```

If you see an error such as “error while loading shared libraries: `libpython3.8.so.1.0`: cannot open shared object file: No such file or directory”, then run the following command:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/local/lib/
```

6. Change the permission of the Python 3.8 installation directory to `755` so that Hue and its related services can leverage the binary and the site packages, as follows:

```
chmod -R 755 /usr/local/lib/python3.8
```

### What to do next

(For Hue) If you have installed Python 3.8 at a custom location, then you must append the custom path in Cloudera Manager Clusters Hue Configuration Hue Service Environment Advanced Configuration Snippet (Safety Valve) separated by colon (`:`) as follows and restart the Hue service:

Key: PATH

Value: `[***CUSTOM-INSTALL-PATH***]:/usr/local/sbin:/usr/local/bin:/usr/sbin:`

### Installing Python 3.8 on SLES 12 for Hue

You must install Python 3.8 on all hosts on which you want to run the Hue service after installing Cloudera Manager and before adding the services to your cluster.

### Before you begin

Install the necessary developer tools such as `gcc` and `make` on your system.

Install the `openssl` package, and place its binaries in the `/lib/` directory. This is needed for installing the MySQL client on MySQL and MariaDB databases.

1. Download the `openssl` package on the host on which you want to install the Python 3.8 package and decompress the file:

```
cd /opt wget https://www.openssl.org/source/openssl-1.0.2o.tar.gz
```

```
tar -xzf openssl-1.0.2o.tar.gz
```

2. Change directory to `openssl-1.0.2o`:

```
cd cd openssl-1.0.2o
```

3. Run the following commands to build and compile the files:

```
./config --shared
make
make install
```

4. Copy the following files to the /lib/ directory:

```
cp libcrypto.so /lib/
cp libssl.so /lib/
cp libcrypto.a /lib/
cp libssl.a /lib/
cp libcrypto.pc /lib/
cp libssl.pc /lib/
cp openssl.pc /lib/
```

5. Delete the openssl-1.0.2o.tar.gz file that you had downloaded:

```
rm ./openssl-1.0.2o.tar.gz
```

6. Configure openssl as follows:

```
custom_openssl = " --with-openssl=/opt/openssl-1.0.2o"
./configure --enable-shared[***CUSTOM-OPENSSL***]
```

Install the following packages before installing Python 3.8:

- libffi-devel-gcc5
- libbz2-devel
- libzip2
- libffi-devel
- libz1
- zlib-devel

```
zypper install -y libffi-devel-gcc5 libbz2-devel libzip2 libffi-devel libz1
zlib-devel
```

## Procedure

1. SSH into the host system as a root user.
2. Download Python 3.8 and decompress the package by running the following commands:

```
cd /opt
curl -O https://www.python.org/ftp/python/3.8.12/Python-3.8.12.tgz
```

```
tar -zxvf Python-3.8.12.tgz
```

3. Change directory to where you decompressed the Python 3.8 package:

```
cd /opt/Python-3.8.12
```

#### 4. Install Python 3.8 as follows:

```
./configure --enable-shared --prefix=[***CUSTOM-INSTALL-PATH***]
```



**Attention:** By default, Python could be installed in any one of the following locations:

- /usr/bin
- /usr/local/python38/bin
- /usr/local/bin
- /opt/rh/rh-python38/root/usr/bin

If you are installing Python 3.8 in any other location, then you must specify the path using the `--prefix` option. You must also create a symbolic link pointing to `/usr/bin` or to `/usr/local/bin` by running the following command:

```
ln -s [***SOURCE***] [***DESTINATION***]
```

For example:

```
ln -s [***CUSTOM-PYTHON-PATH***]/python3.8 /usr/bin/python3.8
```

The `--enable-shared` option is used to build a shared library instead of a static library.

#### 5. Build Python 3.8 as follows:

- a) Run the `make` command to compile the files:

```
make
```

- b) Run the following command to put the compiled files in the default location or in the custom location that you specified using the `--prefix` option:

```
make install
```

- c) Copy the shared compiled library files (`libpython3.8.so`) to the `/lib64/` directory:

```
cp --no-clobber ./libpython3.8.so* /lib64/
```

The `--no-clobber` option is used to prevent overwriting files.

- d) Change the permissions of the `libpython3.8.so` files as follows:

```
chmod 755 /lib64/libpython3.8.so*
```

If you see an error such as “error while loading shared libraries: `libpython3.8.so.1.0`: cannot open shared object file: No such file or directory”, then run the following command:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/local/lib/
```

6. Change the permission of the Python 3.8 installation directory to 755 so that Hue and its related services can leverage the binary and the site packages, as follows:

```
chmod -R 755 /usr/local/lib/python3.8
```

### What to do next

- (For Hue) If you have installed Python 3.8 at a custom location, then you must append the custom path in Cloudera Manager Clusters Hue Configuration Hue Service Environment Advanced Configuration Snippet (Safety Valve) separated by colon (:) as follows and restart the Hue service:

Key: PATH

Value: [\*\*\*CUSTOM-INSTALL-PATH\*\*\*]:usr/local/sbin:usr/local/bin:usr/sbin:

- Clean up the compiled openssl artifacts by running the following commands:

```
rm /lib/libcrypto.so
rm /lib/libssl.so
rm /lib/libcrypto.a
rm /lib/libssl.a
rm /lib/libcrypto.pc
rm /lib/libssl.pc
rm /lib/openssl.pc
```

### Installing Python 3.8 on Ubuntu 18 for Hue

You must install Python 3.8 on all hosts on which you want to run the Hue service after installing Cloudera Manager and before adding the services to your cluster.

### Before you begin

Install the necessary developer tools such as gcc and make on your system.

Install the following packages before installing Python 3.8:

- build-essential
- zlib1g-dev
- libssl-dev
- libffi-dev
- libbz2-dev

```
sudo apt install -y build-essential zlib1g-dev libssl-dev libffi-dev libbz2-dev
```

### Procedure

1. SSH into the host system as a root user.
2. Download Python 3.8 and decompress the package by running the following commands:

```
cd /opt
curl -O https://www.python.org/ftp/python/3.8.12/Python-3.8.12.tgz
```

```
tar -zxvf Python-3.8.12.tgz
```

3. Change directory to where you decompressed the Python 3.8 package:

```
cd /opt/Python-3.8.12
```

#### 4. Install Python 3.8 as follows:

```
./configure --enable-shared --prefix=[***CUSTOM-INSTALL-PATH***]
```



**Attention:** By default, Python could be installed in any one of the following locations:

- /usr/bin
- /usr/local/python38/bin
- /usr/local/bin
- /opt/rh/rh-python38/root/usr/bin

If you are installing Python 3.8 in any other location, then you must specify the path using the `--prefix` option. You must also create a symbolic link pointing to `/usr/bin` or to `/usr/local/bin` by running the following command:

```
ln -s [***SOURCE***] [***DESTINATION***]
```

For example:

```
ln -s [***CUSTOM-PYTON-PATH***]/python3.8 /usr/bin/python3.8
```

The `--enable-shared` option is used to build a shared library instead of a static library.

#### 5. Build Python 3.8 as follows:

- a) Run the `make` command to compile the files:

```
make
```

- b) Run the following command to put the compiled files in the default location or in the custom location that you specified using the `--prefix` option:

```
make install
```

- c) Copy the shared compiled library files (`libpython3.8.so`) to the `/lib64/` directory:

```
cp --no-clobber ./libpython3.8.so* /lib64/
```

The `--no-clobber` option is used to prevent overwriting files.

- d) Change the permissions of the `libpython3.8.so` files as follows:

```
chmod 755 /lib64/libpython3.8.so*
```

If you see an error such as “error while loading shared libraries: `libpython3.8.so.1.0`: cannot open shared object file: No such file or directory”, then run the following command:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/local/lib/
```

6. Change the permission of the Python 3.8 installation directory to `755` so that Hue and its related services can leverage the binary and the site packages, as follows:

```
chmod -R 755 /usr/local/lib/python3.8
```

#### What to do next

(For Hue) If you have installed Python 3.8 at a custom location, then you must append the custom path in Cloudera Manager Clusters Hue Configuration Hue Service Environment Advanced Configuration Snippet (Safety Valve) separated by colon (`:`) as follows and restart the Hue service:

Key: PATH

Value: [\*\*\*CUSTOM-INSTALL-PATH\*\*\*]/usr/local/sbin:/usr/local/bin:/usr/sbin:

### Installing the psycopg2 Python package for PostgreSQL database

If you are using PostgreSQL as a backend database for Hue on 7.x.x, then you must install a version of the psycopg2 package to be at least 2.9.5 on all Hue hosts. The psycopg2 package is automatically installed as a dependency of Agent, but the version installed is often lower than 2.9.3.

Before you begin, you must disable the postgresql10 section from the cloudera-manager.repo file as follows:

1. SSH in to the host as an Administrator.
2. Change to the directory where you had downloaded the cloudera-manager.repo file. On RHEL, the file is present under the /etc/yum.repos.d directory.
3. Open the file for editing and update the value of the enabled property to 0 as follows:

```
[postgresql10]
name=Postgresql 10
baseurl=https://archive.cloudera.com/postgresql10/redhat8/
gpgkey=https://archive.cloudera.com/postgresql10/redhat8/RPM-GPG-KEY-PG
DG-10
enabled=0
gpgcheck=1
module_hotfixes=true
```

4. Save the file and exit.



**Note:** The steps to disable the postgresql10 section are applicable to all supported operating systems (CentOS, RHEL, SLES, and Ubuntu).

#### For CentOS RHEL OEL 7 8

The following steps apply to CentOS 7, RHEL 7, OEL 7, CentOS 8, RHEL 8, and OEL 8:

1. SSH into the Hue server host as a root user.
2. Install the psycopg2-binary package as follows:

```
pip3.8 install psycopg2-binary
```

3. Repeat these steps on all the Hue server hosts.

If you get the "Error: pg\_config executable not found" error while installing the psycopg2-binary package, then run the following commands to install the postgresql, postgresql-devel, python-devel packages:

```
yum install postgresql postgresql-devel python-devel
```

#### For RHEL 9

The following steps apply to RHEL 9, as the minimum version of Python is 3.9:

1. SSH into the Hue server host as a root user.
2. Install pip for Python as follows:

```
yum install python3-pip -y
```

3. Add the /usr/local/bin path to the PATH environment variable:

```
export PATH=$PATH:/usr/local/bin
echo $PATH
```

4. Install the psycopg2-binary package as follows:

```
pip3 install psycopg2-binary
```

5. Repeat these steps on all the Hue server hosts.

If you get the "Error: pg\_config executable not found" error while installing the psycopg2-binary package, then run the following commands to install the postgresql, postgresql-devel, python-devel packages:

```
yum install postgresql postgresql-devel python-devel
```

### For SLES

The following steps apply to SLES 12 (upgrades to Cloudera Runtime 7.1.8) and SLES 15 SP4 (upgrades to Cloudera Runtime 7.1.9 or higher):

1. SSH into the Hue host as a root user.
2. Install the psycopg2 package dependencies by running the following commands:

```
zypper install xmlsec1  
zypper install xmlsec1-devel  
zypper install xmlsec1-openssl-devel
```

3. Install the postgresql-devel package corresponding to your database version by running the following command:

```
zypper -n postgresql[***DB-VERSION***]-devel
```

4. Add the location of the installed postgresql-devel package to the PATH environment variable by running the following command:

```
export PATH=$PATH:/usr/local/bin
```

5. Install the psycopg2 package by running the following command:

```
pip3.8 install psycopg2==2.9.3 --ignore-installed
```

### For Ubuntu

The following steps apply to Ubuntu 18 (upgrades to Cloudera Runtime 7.1.8) and Ubuntu 20 (upgrades to Cloudera Runtime 7.1.9 or higher):

1. SSH into the Hue host as a root user.
2. Install the psycopg2 package dependencies by running the following commands:

```
apt-get install -y xmlsec1  
apt-get install libxmlsec1-openssl  
apt-get install libpq-dev python3-pip -y
```

3. Install the python3-dev and libpq-dev packages by running the following command:

```
apt install python3-dev libpq-dev
```

4. Add the location of the installed postgresql-devel package to the PATH environment variable by running the following command:

```
export PATH=$PATH:/usr/local/bin
```

5. Install the psycopg2 package by running the following command:

```
pip3.8 install psycopg2==2.9.3 --ignore-installed
```

### Installing MySQL client for MySQL databases

To use MySQL as a backend database for Hue, you must install the MySQL client and other required dependencies on all the Hue hosts based on your operating system.



**Attention:** The version 2.2.0 of the MySQL client requires that you set the values of the `MYSQLCLIENT_CFLAGS` and `MYSQLCLIENT_LDFLAGS` environment variables, as follows:

```
$ export MYSQLCLIENT_CFLAGS=`pkg-config mysqlclient --cflags`
$ export MYSQLCLIENT_LDFLAGS=`pkg-config mysqlclient --libs`
```

The version 2.2.0 of the MySQL client also requires you to install the Python 3 and MySQL development headers and libraries, as follows on Debian or Ubuntu operating systems:

```
sudo apt-get install python3-dev default-libmysqlclient-dev build-essential
```

or as follows on CentOS and RHEL operating systems:

```
sudo yum install python3-devel mysql-devel
```

Alternatively, you can install and use the version 2.1.1 of the MySQL client, as follows, which does not have this requirement:

```
pip3 install mysqlclient=2.1.1
```

### For Cent OS

1. SSH into the Hue host as a root user.
2. Download the MySQL yum repository as follows:

```
curl -sSLO https://dev.mysql.com/get/mysql80-community-release-el7-5.noarch.rpm
```

3. Install the package as follows:

```
rpm -ivh mysql80-community-release-el7-5.noarch.rpm
```

4. Install the required dependencies as follows:

```
yum install mysql-devel
yum install -y xmlsec1 xmlsec1-openssl
```

For MySQL version 8.0.27, add the `mysql-community-client-8.0.25` client package as follows:

```
yum install mysql-community-client-8.0.25
```

5. Add the path where you installed the MySQL client and packages to the `PATH` environment variable as follows:

```
export PATH=/usr/local/bin:$PATH
```

6. Install the MySQL client as follows:

```
pip3.8 install mysqlclient
```

### For RHEL

1. SSH into the Hue host as a root user.



2. Download the MySQL yum repository as follows:

(RHEL 7)

```
curl -sSLO https://dev.mysql.com/get/mysql80-community-release-el7-5.noarch.rpm
```

(RHEL 8)

```
curl -sSLO https://dev.mysql.com/get/mysql80-community-release-el8-8.noarch.rpm
```

(RHEL 9)

```
curl -sSLO https://dev.mysql.com/get/mysql80-community-release-el9-4.noarch.rpm
```

3. Install the package as follows:

(RHEL 7)

```
rpm -ivh mysql80-community-release-el7-5.noarch.rpm
```

(RHEL 8)

```
rpm -ivh mysql80-community-release-el8-8.noarch.rpm
```

(RHEL 9)

```
rpm -ivh mysql80-community-release-el9-4.noarch.rpm
```

4. Install the required dependencies as follows:

```
yum install mysql-devel  
yum install -y xmlsec1 xmlsec1-openssl
```

5. Add the path where you installed the MySQL client and packages to the PATH environment variable as follows:

```
export PATH=/usr/local/bin:$PATH
```

6. Install the MySQL client as follows:

(RHEL 8)

```
pip3.8 install mysqlclient
```

(RHEL 9)

```
pip3.9 install mysqlclient
```

### For SLES

1. SSH into the Hue host as a root user.
2. Install the required packages and dependencies as follows:

```
zypper install libmysqlclient-devel  
zypper install xmlsec1  
zypper install xmlsec1-devel  
zypper install xmlsec1-openssl-devel
```

3. Add the path where you installed the packages to the PATH environment variable as follows:

```
export PATH=/usr/local/bin:$PATH
```

4. Install the MySQL client as follows:

```
pip3.8 install mysqlclient
```

#### For Ubuntu

1. SSH into the Hue host as a root user.
2. Install the required packages and dependencies as follows:

```
apt-get install libmysqlclient-dev  
apt-get install -y xmlsec1  
apt-get install libxmlsec1-openssl
```

3. Add the path where you installed the packages to the PATH environment variable as follows:

```
export PATH=/usr/local/bin:$PATH
```

4. Install the MySQL client as follows:

```
pip3.8 install mysqlclient
```

### Installing MySQL client for MariaDB databases

To use MariaDB as a backend database for Hue, you must install the MySQL client and other required dependencies on all the Hue hosts based on your operating system.

#### For Cent OS

1. SSH into the Hue host as a root user.
2. Install the required dependencies as follows:

```
yum install -y xmlsec1 xmlsec1-openssl
```

3. Add the path where you installed the packages to the PATH environment variable as follows:

```
export PATH=/usr/local/bin:$PATH
```

4. Install the MySQL client as follows:

```
pip3.8 install mysqlclient
```

#### For RHEL

1. SSH into the Hue host as a root user.
2. Install the required dependencies as follows:

```
yum install mysql-devel  
yum install -y xmlsec1 xmlsec1-openssl
```

3. Add the path where you installed the packages to the PATH environment variable as follows:

```
export PATH=/usr/local/bin:$PATH
```

4. Install the MySQL client as follows:

(RHEL 8)

```
pip3.8 install mysqlclient
```

(RHEL 9)

```
pip3.9 install mysqlclient
```

#### For SLES

1. SSH into the Hue host as a root user.
2. Install the required packages and dependencies as follows:

```
zypper install libmysqlclient-devel  
zypper install xmlsec1  
zypper install xmlsec1-devel  
zypper install xmlsec1-openssl-devel
```

3. Add the path where you installed the packages to the PATH environment variable as follows:

```
export PATH=/usr/local/bin:$PATH
```

4. Install the MySQL client as follows:

```
pip3.8 install mysqlclient
```

#### For Ubuntu

1. SSH into the Hue host as a root user.
2. Install the required packages and dependencies as follows:

```
apt-get install libmysqlclient-dev  
apt-get install -y xmlsec1  
apt-get install libxmlsec1-openssl
```

3. Add the path where you installed the packages to the PATH environment variable as follows:

```
export PATH=/usr/local/bin:$PATH
```

4. Install the MySQL client as follows:

```
pip3.8 install mysqlclient
```

## Oozie

Perform the following post migration tasks by validating the database URL, installing the new shared libraries, accessing Oozie Loadbalancer URL, and configuring load balancer.

### Validate Database URL

Validate the database URL.

### Before you begin

Validate the configured Oozie JDBC URL. Before starting the Oozie service, check the value.

### Procedure

1. Verify JDBC Database settings in Oozie service.

- Configure Oozie Load Balancer Hostname, Oozie Load Balancer HTTP Port and Oozie Load Balancer HTTPS Port if OOZIE HA is enabled.

- In Oozie service configurations under Oozie Server Advanced Configuration Snippet (Safety Valve) for oozie-site.xml add oozie.service.AuthorizationService.admin.users configuration with value of oozie, oozie-admin.



**Note:** Replace Oozie with the service account username if it is different from default one.

### Installing the new Shared Libraries

Install new Shared Libraries

#### About this task

Follow the steps:

#### Procedure

- Start Oozie services.
- Set ShareLib Root Directory=`/user/<oozie user>` in Oozie configs
- From Cloudera Manager, navigate to Cloudera Manager > Oozie > Install Oozie Shared Lib



**Note:** Tez libraries are removed from Oozie ShareLib. If you are using Tez on Oozie, then you must manually copy the Tez jar files and execute a ShareLib update.

#### Update Oozie properties

You must manually check and update the properties for Oozie in Cloudera Manager if the same properties had custom values in Ambari. Similarly, check the admin user's configuration.

Manually update the following properties:

- Oozie Server Plugins
- Oozie SchemaService Workflow Extension Schemas
- Oozie ActionService Executor Extension Classes
- Oozie Credential Classes

#### Access Oozie load balancer URL

This is a workaround to access Oozie load balancer URL.

#### Procedure

- Log in to load balancer host.
- In the `/etc/httpd/conf/httpd.conf` file, change `SSLProxyVerify` required to `SSLProxyVerify off`.
- stop httpd `"apachectl -k stop"`.
- Start httpd `"apachectl -k start"`.  
You can now access the Oozie load balancer URL.

## Oozie Load Balancer configuration

To enable Oozie High Availability, you must manually configure a Load Balancer.

### About this task

Cloudera recommends using the HAProxy Load Balancer. These steps explain how to configure the HAProxy load balancer. However, you can choose to configure a different Load Balancer.

### Procedure

1. Install HAProxy on the host where you are setting up and configuring the Oozie load balancer. For more information, see the [HAProxy documentation](#).
2. You must configure the Oozie load balancer for both HTTP and HTTPS ports.

```
This is an example:
global
    log                  127.0.0.1 local2
    pidfile              /var/run/haproxy.pid
    maxconn              4000
    user                 haproxy
    group                haproxy
    daemon
    stats                socket /tmp/haproxy

defaults
    mode                 http
    log                  global
    option                httplog
    option                dontlognull
    option forwardfor    except 127.0.0.0/8
    option                redispatch
    retries               3
    timeout http-request 10s
    timeout queue         1m
    timeout connect      10s
    timeout client        10m
    timeout server        10m
    timeout check         10s
    maxconn              3000

listen admin
    bind *:8000
    stats enable

#-----
# main frontend which proxys to the backends
#-----
frontend                oozie_front
    bind                  *:5000 ssl crt /var/lib/cloudera-scm-agent
    /agent-cert/cdep-host_key_cert_chain_decrypted.pem
    default_backend       oozie

#-----
# round robin balancing between the various backends
#-----
backend oozie
    balance                roundrobin
    server oozie1 my-oozie-host-1:11443/oozie check ssl ca-file /var/lib/
    cloudera-scm-agent/agent-cert/cm-auto-global_cacerts.pem
    server oozie2 my-oozie-host-2:11443/oozie check ssl ca-file /var/lib/
    cloudera-scm-agent/agent-cert/cm-auto-global_cacerts.pem
    server oozie3 my-oozie-host-3:11443/oozie check ssl ca-file /var/lib/
    cloudera-scm-agent/agent-cert/cm-auto-global_cacerts.pem
```

```

#-----
# main frontend which proxys to the http backends
#-----
frontend                                oozie_front_http
    bind                                 *:5002
    default_backend                       oozie_http

#-----
# round robin balancing between the various http backends
#-----
backend oozie_http
    balance                               roundrobin
    server oozie_http1 my-oozie-host-1:11000/oozie check
    server oozie_http2 my-oozie-host-2:11000/oozie check
    server oozie_http3 my-oozie-host-3:11000/oozie check

```

Using the example, the load balancer is setup for three Oozie instances. The load balancer listens on port 5002 for HTTP connections and forwards it to Oozie's port 11000. The load balancer listens on port 5000 for HTTPS connections and forwards it to Oozie's port 11443.

If you not enabled SSL in Oozie, then you do not need the HTTPS load balancer. For HTTPS load balancing, ensure that you set up the certificate.

3. Continue to configure the load balancer by enabling Oozie High Availability. For information about enabling Oozie High Availability, see [Enabling Oozie High Availability](#) .

### Atlas advanced configuration snippet (Safety valve)

If non-default user names are used, then you must set this parameter. This setting is applicable only if you are upgrading from HDP 2.6.5 to CDP Private Cloud Base 7.1.6. This section is not applicable if you are upgrading to CDP Private Cloud Base 7.1.7.

1. Log into Cloudera Manager UI.
2. Navigate to Clusters
3. Select the Atlas service
4. Go to Configurations
5. Search for Atlas Server Environment Advanced Configuration Snippet
6. Rename ATLAS\_OPTS to ATLAS\_CUSTOM\_OPTS

### Migrating Atlas data

After Atlas is available on the Cloudera Private Cloud Base cluster, you must import the Atlas data that the Atlas migration exporter utility exported from the HDP 2.6.5.x cluster.

#### Before you begin

Use these properties to improve the speed of Atlas data import:

For a node with 4 cores and 8 GB of heap space, the estimated duration for import is 0.75 million entities per hour.

1. Go to Atlas Configuration Category Advanced Custom application-properties :
2. atlas.migration.mode.batch.size: The recommended value is 300.
3. atlas.migration.mode.workers: The recommended value is 3 when using one Atlas Solr shard.

The value to be set depends on the following calculation:

if `ATLAS_SOLR_SHARDS > 1`,

`atlas.patch.numWorkers=ATLAS_SOLR_SHARDS * 3`

Additional patches are applied after the migration is completed. These are the properties that will help with improving the speed of patches.

## Procedure

1. Configure Atlas in Cloudera Private Cloud Base with the location of the exported data.
2. Configure the `atlas.migration.data.filename` property.
3. In Cloudera Manager, go to Clusters and select Atlas.
4. From Atlas configuration, set the Advanced Configuration Snippet (Safety Valve) for `conf/atlas-application.properties` value to the location which contains exported Atlas data.
5. Set additional properties:
  - `atlas.migration.mode.batch.size=300`.
  - `atlas.migration.mode.workers=<use the value from calculation above>`
  - `atlas.patch.batchSize=300`
  - `atlas.patch.numWorkers=<use the value from calculation above>`

For example:

- `atlas.migration.data.filename=/var/lib/atlas-data`
- `atlas.migration.mode.batch.size=300`.
- `atlas.migration.mode.workers=3`
- `atlas.patch.batchSize=300`
- `atlas.patch.numWorkers=3`

Atlas Server Advanced Configuration Snippet (Safety Valve) for `conf/atlas-application.properties`

[conf/atlas-application.properties\\_role\\_safety\\_valve](#)

Atlas Server Default Group [Undo](#)

```
atlas.migration.data.filename=/tmp/atlas-data
atlas.migration.mode.batch.size=300
atlas.migration.mode.workers=14
atlas.patch.batchSize=300
atlas.patch.numWorkers=3
```

1 - 1 of 1

6. Save the configuration.
7. Restart Atlas from Cloudera Manager. Go to Atlas Actions Restart .

Cluster 1

ATLAS-1 [Actions](#)

- Start
- Restart **Restart a service.**
- Rolling Restart
- Initialize Atlas
- Stop
- Add Role Instances
- Rename
- Delete
- Enter Maintenance Mode
- Create Ranger Atlas Plugin Audit Directory
- Create HBase tables for Atlas
- Create Kafka Topics for Atlas
- Deploy Client Configuration

1 Edited Value Reason for change: Modified Atlas Server Advanced Configuration Snippet (Safety Valve) for conf/atlas-application.properties

[Save Changes\(ctrl+S\)](#)

8. Atlas starts in the migration mode and data import should commence. During the migration process, Atlas blocks all the REST API calls and Atlas Hook notification processing.

9. To check the migration status, use the following command:

```
http://[atlas_server]:21000/api/atlas/admin/status
The migration status is displayed in the browser window:
{"Status": "Migration", "currentIndex": 139, "percent": 67, "
startTimeUTC": "2018-04-06T00:54:53.399Z" }
```

10. The progress of import can be monitored using Atlas logs on the node where the migration is running. On the completion of migration, the log should have this entry: Done! loadLegacyGraphSON. (GraphDBGGraphSONMigrator:76)
11. After the migration is complete, change the status of Atlas from migration mode to normal operation by removing the atlas.migration.data.filename property and restarting Atlas in Cloudera Manager.



**Note:**

- After Atlas is up and running, it continues processing messages from the integrated services. Explicitly importing data through import scripts is not required.
- As part of this migration, the value of atlas.cluster.name is mapped to the new property atlas.metadata.namespace. This value is used within Atlas for generating unique names of entities.
- You must restart Atlas and then remove the property atlas.migration.mode.hdp.to.cdp=true from the Atlas Server Advanced Configuration Snippet (Safety Valve) for conf/atlas-application.properties after the Solr collections are migrated to the CDP cluster.




## Phoenix

Add the Apache Phoenix service. If you are using Phoenix Query Server (PQS) in your source HDP deployment, you must manually add the Apache Phoenix service using Cloudera Manager to complete the Phoenix upgrade.

### About this task

You must add the Apache Phoenix service using Cloudera Manager. To add the Apache Phoenix service:

### Procedure

- In Cloudera Manager, click Home .
-  Click  to the right of the cluster name and select Add Service. A list of service types display. Select Phoenix from the list and click Continue. Follow the wizard to add the Phoenix service.

### Map Phoenix schemas to HBase namespaces

HBase namespaces enable tighter control of where a particular data set is stored on the HBase RegionServers. If you have mapped Phoenix schemas to HBase namespaces in your source HDP cluster, you must enable namespace mapping by configuring a set of properties using Cloudera Manager.

You can check if you have enabled namespace mapping in source HDP deployment Ambari's HBase service. Check the following HBase configurations in your source HDP deployment to see if you have mapped Phoenix schemas to HBase namespaces.

| HBase Config                              | Value |
|---|-------|
| phoenix.schema.isNamespaceMappingEnabled  | true  |
| phoenix.schema.mapSystemTablesToNamespace | true  |

If you have enabled namespace mapping in HDP, and want to use in CDP, you must enable namespace mapping in CDP using Cloudera Manager. For more information, see *Enable namespace mapping* using the link in the related information section.



## Related Information

[Enable namespace mapping](#)

## Starting all services

You must now start all the services.

Start all the services and ensure that all the services of Cloudera Private Cloud Base are up and running.

## Hive Policy Additions

Hive Metastore Canary test might fail on clusters migrated from HDP because the service account used to perform these tests is different between the two management consoles.

The new service account is Hue. You must add the Hue user to the policies defined that hold Hive as the superuser. If there is no policy in Ranger defining Hive as a superuser across key directories in HDFS, then you are using impersonation. Impersonation is no longer the standard or recommended. Consult the Cloudera team for details.

Sample Hive or Hue Privileged User Policy

The screenshot shows the 'Edit Policy' page in Cloudera Manager. The breadcrumb trail is 'Service Manager > cm\_hdfs Policies > Edit Policy'. The policy details are as follows:

- Policy Type:** Access
- Policy ID:** 282
- Policy Name:** HDP BP - Hive Superuser Access. Status: enabled (radio button selected), normal (radio button unselected).
- Policy Label:** Policy Label
- Resource Path:** A list of paths: /hdp, /apps/hive/warehouse, /apps/spark/warehouse, /warehouse/tablespace/managed, /warehouse/tablespace/external, /tmp, and /user/tez. A 'recursive' radio button is selected.
- Description:** Best Practices HDFS Permissions for Hive Superuser doas=false
- Audit Logging:** YES (radio button selected)

Below the policy details is the 'Allow Conditions' section, which is currently hidden. It contains a table with the following columns: Select Role, Select Group, Select User, Permissions, Delegate Admin, and a red 'X' button.

| Select Role  | Select Group  | Select User          | Permissions                | Delegate Admin |   |
|--------------|---------------|----------------------|----------------------------|----------------|---|
| Select Roles | Select Groups | hive, ambari-qa, hue | Read, Write, Execute, Edit |                | X |

## Knox

The following post migration steps are optional.

### About this task

After migrating to CDP Private Cloud Base, you can use the cdp-proxy (UIs) and cdp-proxy-api (APIs) topologies to access services through Knox. To use advanced gateway, some of the following post-migration steps are required.



**Note:** In CDP Private Cloud Base, SSO is provided to the platform through the proxying of UIs via cdp-proxy topology.

### Topology migration

Migrate the custom HDP topologies to the CDP cluster.

#### Before you begin

There are existing Knox topologies in the CDP cluster available. If you have any custom topologies on the HDP cluster and you want to migrate to the CDP cluster, then perform the following steps:

#### Procedure

1. Add a provider configuration with the details from the providers section of the topology (or provider configuration) to be migrated. If the provider configuration contains any credential aliases, then you must add them. For more information on adding a provider configuration, see [Add a provider configuration](#) and for aliases, see [Saving Aliases](#)
2. Add a descriptor with the services and corresponding parameters from the topology (or descriptor) to be migrated, referencing the newly added provider configuration. For more information on adding a descriptor, see [Add a descriptor](#)

### Migrate Credential Aliases

Migrate the HDP credential aliases to the CDP cluster.

#### Before you begin

If you want to migrate Credential Aliases, then see [Saving Aliases](#)

### Migrate signing key

Migrate Knox's signing key to the CDP cluster.

#### Before you begin

To add Knox signing key, you must add the credential aliases and configure the gateway properties.

#### Procedure

1. If you want to migrate the signing key then the following credential aliases must be added with the corresponding values from the HDP signing key store:
  - gateway.signing.keystore.password.alias - The password for the signing keystore.
  - gateway.signing.key.passphrase.alias - The passphrase to sign the private key stored in the signing keystore.
2. In addition, the following gateway properties must be configured with the Cloudera Manager UI:
  - gateway.signing.keystore.name
  - gateway.signing.key.alias
  - gateway.signing.keystore.type
  - gateway.signing.keystore.password.alias
  - gateway.signing.key.passphrase.alias

### Configure Apache Knox authentication for AD/LDAP

Knox authentication configurations for LDAP and AD in Cloudera Manager.

#### About this task

For more information on configuring Knox LDAP in Cloudera Manager, see [Configuring Knox LDAP](#)

## Client Configurations

Migrating from Ambari to Cloudera Manager can leave the Ambari-managed HDP artifacts and links that may not have changed. After everything has been configured, deploy Client Configuration [Cloudera Manager > Clusters > Actions > Deploy Client Configurations]. This fixes any missing configuration references in the cluster.

## Securing ZooKeeper

By default, the AM2CM tool adds the `-Dzookeeper.skipACL=yes` configuration to assist with the migration. You must remove the `-Dzookeeper.skipACL=yes` configuration under Java Configuration Options for Zookeeper Service to secure ZooKeeper and restart the service.

## Zeppelin Shiro configurations

If you had Zeppelin Shiro configurations in the HDP cluster, then you must configure them manually on the CDP Private Cloud Base cluster.

### About this task

Perform the following steps:



**Important: Deprecation notice for Zeppelin:** Zeppelin will be deprecated in Cloudera Runtime 7.1.9. For more information, see *Deprecation Notices in Cloudera Runtime 7.1.9*.

### Procedure

1. Log into Cloudera Manager
2. Navigate to Clusters
3. Select the Zeppelin service
4. Click the Configuration tab
5. Search for Shiro and add the configurations
6. Click Save Changes

## Migrating Spark workloads to

Migrating Spark workloads from CDH or HDP to involves learning the Spark semantic changes in your source cluster and the target cluster. You get details about how to handle these changes.

### Spark 1.6 to Spark 2.4 Refactoring

Because Spark 1.6 is not supported on , you need to refactor Spark workloads from Spark 1.6 on CDH or HDP to Spark 2.4 on .

This document helps in accelerating the migration process, provides guidance to refactor Spark workloads and lists migration. Use this document when the platform is migrated from CDH or HDP to .

### Handling prerequisites

You must perform a number of tasks before refactoring workloads.

### About this task

Assuming all workloads are in working condition, you perform this task to meet refactoring prerequisites.

### Procedure

1. Identify all the workloads in the cluster (CDH/HDP) which are running on Spark 1.6 - 2.3.

## 2. Classify the workloads.

Classification of workloads will help in clean-up of the unwanted workloads, plan resources and efforts for workload migration and post upgrade testing.

Example workload classifications:

- Spark Core (scala)
- Java-based Spark jobs
- SQL, Datasets, and DataFrame
- Structured Streaming
- MLlib ()
- PySpark (Python on Spark)
- Batch Jobs
- Scheduled Jobs
- Ad-Hoc Jobs
- Critical/Priority Jobs
- Huge data Processing Jobs
- Time taking jobs
- Resource Consuming Jobs etc.
- Failed Jobs

Identify configuration changes

## 3. Check the current Spark jobs configuration.

- Spark 1.6 - 2.3 workload configurations which have dependencies on job properties like scheduler, old python packages, classpath jars and might not be compatible post migration.
- In , Capacity Scheduler is the default and recommended scheduler. Follow [Fair Scheduler to Capacity Scheduler transition](#) guide to have all the required queues configured in the cluster post upgrade. If any configuration changes are required, modify the code as per the new capacity scheduler configurations.
- For workload configurations, see the Spark History server UI [http://spark\\_history\\_server:18088/history/<application\\_number>/environment/](http://spark_history_server:18088/history/<application_number>/environment/).

## 4. Identify and capture workloads having data storage locations (local and HDFS) to refactor the workloads post migration.

## 5. Refer to [unsupported Apache Spark features](#), and plan refactoring accordingly.

### Spark 1.6 to Spark 2.4 changes

A description of the change, the type of change, and the required refactoring provide the information you need for migrating from Spark 1.6 to Spark 2.4.

#### *New Spark entry point SparkSession*

There is a new Spark API entry point: SparkSession.

Type of change

Syntactic/Spark core

Spark 1.6

Hive Context and SQLContext, such as import SparkContext, HiveContext are supported.

Spark 2.4

SparkSession is now the entry point.

Action Required

Replace the old SQLContext and HiveContext with SparkSession. For example:

```
import org.apache.spark.sql.SparkSession
val spark = SparkSession
  .builder()
  .appName("Spark SQL basic example")
```

```
.config("spark.some.config.option", "some-value")
.getOrCreate()
```

### *Dataframe API registerTempTable deprecated*

The Dataframe API registerTempTable has been deprecated in Spark 2.4.

Type of change:

Syntactic/Spark core change

Spark 1.6

registerTempTable is used to create a temporary table on a Spark dataframe. For example, df.registerTempTable('tmpTable').

Spark 2.4

registerTempTable is deprecated.

Action Required

Replace registerTempTable using createOrReplaceTempView. df.createOrReplaceTempView('tmpTable').

### *union replaces unionAll*

The dataset and DataFrame API unionAll has been deprecated and replaced by union.

Type of change: Syntactic/Spark core change

Spark 1.6

unionAll is supported.

Spark 2.4

unionAll is deprecated and replaced by union.

Action Required

Replace unionAll with union. For example: val df3 = df.unionAll(df2) with val df3 = df.union(df2)

### *Empty schema not supported*

Writing a dataframe with an empty or nested empty schema using any file format, such as parquet, orc, json, text, or csv is not allowed.

Type of change: Syntactic/Spark core

Spark 1.6 - 2.3

Writing a dataframe with an empty or nested empty schema using any file format is allowed and will not throw an exception.

Spark 2.4

An exception is thrown when you attempt to write dataframes with empty schema. For example, if there are statements such as df.write.format("parquet").mode("overwrite").save(somePath), the following error occurs: org.apache.spark.sql.AnalysisException: Parquet data source does not support null data type.

Action Required

Make sure that DataFrame is not empty. Check whether DataFrame is empty or not as follows:

```
if (!df.isEmpty) df.write.format("parquet").mode("overwrite").save("somePath")
```

### *Referencing a corrupt JSON/CSV record*

In Spark 2.4, queries from raw JSON/CSV files are disallowed when the referenced columns only include the internal corrupt record column.

Type of change: Syntactic/Spark core

Spark 1.6

A query can reference a `_corrupt_record` column in raw JSON/CSV files.

Spark 2.4

An exception is thrown if the query is referencing `_corrupt_record` column in these files. For example, the following query is not allowed: `spark.read.schema(schema).json(file).filter($"_corrupt_record".isNotNull).count()`

Action Required

Cache or save the parsed results, and then resend the query.

```
val df = spark.read.schema(schema).json(file).cache()
df.filter($"_corrupt_record".isNotNull).count()
```

*Dataset and DataFrame API explode deprecated*

Dataset and DataFrame API explode has been deprecated.

Type of change: Syntactic/Spark SQL change

Spark 1.6

Dataset and DataFrame API explode are supported.

Spark 2.4

Dataset and DataFrame API explode have been deprecated. If explode is used, for example `dataframe.explode()`, the following warning is thrown:

```
warning: method explode in class Dataset is deprecated: use flatMap() or select() with functions.explode() instead
```

Action Required

Use `functions.explode()` or `flatMap` (`import org.apache.spark.sql.functions.explode`).

*CSV header and schema match*

Column names of csv headers must match the schema.

Type of change: Configuration/Spark core changes

Spark 1.6 - 2.3

Column names of headers in CSV files are not checked against the against the schema of CSV data.

Spark 2.4

If columns in the CSV header and the schema have different ordering, the following exception is thrown:`java.lang.IllegalArgumentException: CSV file header does not contain the expected fields.`

Action Required

Make the schema and header order match or set `enforceSchema` to false to prevent getting an exception. For example, read a file or directory of files in CSV format into Spark DataFrame as follows: `df3 = spark.read.option("delimiter", ";").option("header", True).option("enforeSchema", False).csv(path)`

The default "header" option is true and `enforceSchema` is False.

If `enforceSchema` is set to true, the specified or inferred schema will be forcibly applied to datasource files, and headers in CSV files are ignored. If `enforceSchema` is set to false, the schema is validated against all headers in CSV files when the header option is set to true. Field names in the schema and column names in CSV headers are checked by their positions taking into account `spark.sql.caseSensitive`. Although the default value is true, you should disable the `enforceSchema` option to prevent incorrect results.

### Table properties support

Table properties are taken into consideration while creating the table.

Type of change: Configuration/Spark Core Changes

Spark 1.6 - 2.3

Parquet and ORC Hive tables are converted to Parquet or ORC by default, but table properties are ignored. For example, the compression table property is ignored:

```
CREATE TABLE t(id int) STORED AS PARQUET TBLPROPERTIES (parquet.compression
'NONE')
```

This command generates Snappy Parquet files.

Spark 2.4

Table properties are supported. For example, if no compression is required, set the TBLPROPERTIES as follows: (parquet.compression 'NONE').

This command generates uncompressed Parquet files.

Action Required

Check and set the desired TBLPROPERTIES.

### CREATE OR REPLACE VIEW and ALTER VIEW not supported

ALTER VIEW and CREATE OR REPLACE VIEW AS commands are no longer supported.

Type of change: Configuration/Spark Core Changes

Spark 1.6

You can create views as follows:

```
CREATE OR REPLACE [ [ GLOBAL ] TEMPORARY ] VIEW [ IF NOT EXISTS ] view_name
[ column_list ]
[ COMMENT view_comment ]
[ properties ]
AS query

ALTER VIEW view_name
{ rename |
set_properties |
unset_properties |
alter_body }
```

Spark 2.4

ALTER VIEW and CREATE OR REPLACE commands above are not supported.

Action Required

Recreate views using the following syntax:

```
CREATE [ [ GLOBAL ] TEMPORARY ] VIEW [ IF NOT EXISTS ] view_name
[ column_list ]
[ COMMENT view_comment ]
[ properties ]
AS query
```

### Managed table location

Creating a managed table with nonempty location is not allowed.

Type of change: Property/Spark core changes

Spark 1.6 - 2.3

You can create a managed table having a nonempty location.

#### Spark 2.4

Creating a managed table with nonempty location is not allowed. In Spark 2.4, an error occurs when there is a write operation, such as `df.write.mode(SaveMode.Overwrite).saveAsTable("testdb.testtable")`. The error side-effects are the cluster is terminated while the write is in progress, a temporary network issue occurs, or the job is interrupted.

#### Action Required

Set `spark.sql.legacy.allowCreatingManagedTableUsingNonemptyLocation` to true at runtime as follows:

```
spark.conf.set("spark.sql.legacy.allowCreatingManagedTableUsingNonemptyLocation", "true")
```

#### *Write to Hive bucketed tables*

Type of change: Property/Spark SQL changes

#### Spark 1.6

By default, you can write to Hive bucketed tables.

#### Spark 2.4

By default, you cannot write to Hive bucketed tables.

For example, the following code snippet writes the data into a bucketed Hive table:

```
newPartitionsDF.write.mode(SaveMode.Append).format("hive").insertInto(hive_test_db.test_bucketing)
```

The code above will throw the following error:

```
org.apache.spark.sql.AnalysisException: Output Hive table `hive_test_db`.`test_bucketing` is bucketed but Spark currently does NOT populate bucketed output which is compatible with Hive.
```

#### Action Required

To write to a Hive bucketed table, you must use `hive.enforce.bucketing=false` and `hive.enforce.sorting=false` to forego bucketing guarantees.

#### *Rounding in arithmetic operations*

Arithmetic operations between decimals return a rounded value, instead of NULL, if an exact representation is not possible.

Type of change: Property/Spark SQL changes

#### Spark 1.6

Arithmetic operations between decimals return a NULL value if an exact representation is not possible.

#### Spark 2.4

The following changes have been made:

- Updated rules determine the result precision and scale according to the SQL ANSI 2011.
- Rounding of the results occur when the result cannot be exactly represented with the specified precision and scale instead of returning NULL.
- A new config `spark.sql.decimalOperations.allowPrecisionLoss` which default to true (the new behavior) to allow users to switch back to the old behavior. For example, if your code includes import statements that resemble those below, plus arithmetic operations, such as multiplication and addition, operations are performed using dataframes.

```
from pyspark.sql.types import DecimalType
```



```
from decimal import Decimal
```

#### Action Required

If precision and scale are important, and your code can accept a NULL value (if exact representation is not possible due to overflow), then set the following property to false. `spark.sql.decimalOperations.allowPrecisionLoss = false`

#### *Precedence of set operations*

Set operations are executed by priority instead having equal precedence.

Type of change: Property/Spark SQL changes

Spark 1.6 - 2.3

If the order is not specified by parentheses, equal precedence is given to all set operations.

Spark 2.4

If the order is not specified by parentheses, set operations are performed from left to right with the exception that all INTERSECT operations are performed before any UNION, EXCEPT or MINUS operations.

For example, if your code includes set operations, such as INTERSECT , UNION, EXCEPT or MINUS, consider refactoring.

#### Action Required

Change the logic according to following rule:

If the order of set operations is not specified by parentheses, set operations are performed from left to right with the exception that all INTERSECT operations are performed before any UNION, EXCEPT or MINUS operations.

If you want the previous behavior of equal precedence then, set `spark.sql.legacy.setopsPrecedence.enabled=true`.

#### *HAVING without GROUP BY*

HAVING without GROUP BY is treated as a global aggregate.

Type of change: Property/Spark SQL changes

Spark 1.6 - 2,3

HAVING without GROUP BY is treated as WHERE. For example, `SELECT 1 FROM range(10) HAVING true` is executed as `SELECT 1 FROM range(10) WHERE true`, and returns 10 rows.

Spark 2.4

HAVING without GROUP BY is treated as a global aggregate. For example, `SELECT 1 FROM range(10) HAVING true` returns one row, instead of 10, as in the previous version.

#### Action Required

Check the logic where having and group by is used. To restore previous behavior, set `spark.sql.legacy.parser.havingWithoutGroupByAsWhere=true`.

#### *CSV bad record handling*

How Spark treats malformations in CSV files has changed.

Type of change: Property/Spark SQL changes

Spark 1.6 - 2.3

CSV rows are considered malformed if at least one column value in the row is malformed. The CSV parser drops malformed rows in the DROPMALFORMED mode or outputs an error in the FAILFAST mode.

Spark 2.4

A CSV row is considered malformed only when it contains malformed column values requested from CSV datasource, other values are ignored.

#### Action Required

To restore the Spark 1.6 behavior, set `spark.sql.csv.parser.columnPruning.enabled` to `false`.

### Spark 2.4 CSV example

A CSV example illustrates the CSV-handling change in Spark 2.4.

In the following CSV file, the first two records describe the file. These records are not considered during processing and need to be removed from the file. The actual data to be considered for processing has three columns (jersey, name, position).

```
These are extra line1
These are extra line2
10,Messi,CF
7,Ronaldo,LW
9,Benzema,CF
```

The following schema definition for the DataFrame reader uses the option `DROPMALFORMED`. You see only the required data; all the description and error records are removed.

```
schema=Structtype([Structfield("jersey",StringType()),Structfield("name",StringType()),Structfield("position",StringType())])
df1=spark.read\
.option("mode","DROPMALFORMED")\
.option("delimiter",",")\
.schema(schema)\
.csv("inputfile")
df1.select("*").show()
```

Output is:

| jersey | name    | position |
|--------|---------|----------|
| 10     | Messi   | CF       |
| 7      | Ronaldo | LW       |
| 9      | Benzema | CF       |

Select two columns from the dataframe and invoke `show()`:

```
df1.select("jersey","name").show(truncate=False)
```

| jersey                | name    |
|-----------------------|---------|
| These are extra line1 | null    |
| These are extra line2 | null    |
| 10                    | Messi   |
| 7                     | Ronaldo |
| 9                     | Benzema |

Malformed records are not dropped and pushed to the first column and the remaining columns will be replaced with null. This is due to the CSV parser column pruning which is set to true by default in Spark 2.4.

Set the following conf, and run the same code, selecting two fields.

```
spark.conf.set("spark.sql.csv.parser.columnPruning.enabled",False)
```

```
df2=spark.read\
.option("mode","DROPMALFORMED")\
.option("delimiter",",")\
```

```
.schema(schema)\
.csv("inputfile")
df2.select("jersey","name").show(truncate=False)
```

| jersey | name    |
|--------|---------|
| 10     | Messi   |
| 7      | Ronaldo |
| 9      | Benzema |

Conclusion: If working on selective columns, to handle bad records in CSV files, set `spark.sql.csv.parser.columnPruning.enabled` to false; otherwise, the error record is pushed to the first column, and all the remaining columns are treated as nulls.

### Configuring storage locations

To execute the workloads in , you must modify the references to storage locations. In , references must be changed from HDFS to a cloud object store such as S3.

### About this task

The following sample query shows a Spark 2.4 HDFS data location.

```
scala> spark.sql("CREATE TABLE IF NOT EXISTS default.sales_spark_2(Region string, Country string,Item_Type string,Sales_Channel string,Order_Priority string,Order_Date date,Order_ID int,Ship_Date date,Units_sold string,Unit_Price string,Unit_cost string,Total_revenue string,Total_Cost string,Total_Profit string) row format delimited fields terminated by ', '")
scala> spark.sql("load data local inpath '/tmp/sales.csv' into table default.sales_spark_2")
scala> spark.sql("select count(*) from default.sales_spark_2").show()
```

The following sample query shows a Spark 2.4 S3 data location.

```
scala> spark.sql("CREATE TABLE IF NOT EXISTS default.sales_spark_2(Region string, Country string,Item_Type string,Sales_Channel string,Order_Priority string,Order_Date date,Order_ID int,Ship_Date date,Units_sold string,Unit_Price string,Unit_cost string,Total_revenue string,Total_Cost string,Total_Profit string) row format delimited fields terminated by ', '")
scala> spark.sql("load data inpath 's3://<bucket>/sales.csv' into table default.sales_spark_2")
scala> spark.sql("select count(*) from default.sales_spark_2").show()
```

### Querying Hive managed tables from Spark

Hive-on-Spark is not supported on . You need to use the Hive Warehouse Connector (HWC) to query Apache Hive managed tables from Apache Spark.

To read Hive external tables from Spark, you do not need HWC. Spark uses native Spark to read external tables. For more information, see the [Hive Warehouse Connector documentation](#).

The following example shows how to query a Hive table from Spark using HWC:

```
spark-shell --jars /opt/cloudera/parcels/CDH/jars/hive-warehouse-connector-assembly-1.0.0.7.1.4.0-203.jar --conf spark.sql.hive.hiveserver2.jdbc.url=jdbc:hive2://cdhhd02.uddepta-bandyopadhyay-s-account.cloud:10000/default --conf spark.sql.hive.hiveserver2.jdbc.url.principal=hive/cdhhd02.uddepta-bandyopadhyay-s-account.cloud@Uddepta-bandyopadhyay-s-Account.CLOUD
scala> val hive = com.hortonworks.hwc.HiveWarehouseSession.session(spark).build()
scala> hive.executeUpdate("UPDATE hive_acid_demo set value=25 where key=4")
scala> val result=hive.execute("select * from default.hive_acid_demo")
```

```
scala> result.show()
```

### Compiling and running Spark workloads

After modifying the workloads, compile and run (or dry run) the refactored workloads on Spark 2.4.

You can write Spark applications using Java, Scala, Python, SparkR, and others. You build jars from these scripts using one of the following compilers.

- Java (with Maven/Java IDE),
- Scala (with sbt),
- Python (pip).
- SparkR (RStudio)

### Compiling and running a Java-based job

You see by example how to compile a Java-based Spark job using Maven.

### About this task

In this task, you see how to compile the following example Spark program written in Java:

```
/* SimpleApp.java */
import org.apache.spark.sql.SparkSession;
import org.apache.spark.sql.Dataset;

public class SimpleApp {
    public static void main(String[] args) {
        String logFile = "YOUR_SPARK_HOME/README.md"; // Should be some file on
        your system
        SparkSession spark = SparkSession.builder().appName("Simple Applicatio
        n").getOrCreate();
        Dataset<String> logData = spark.read().textFile(logFile).cache();

        long numAs = logData.filter(s -> s.contains("a")).count();
        long numBs = logData.filter(s -> s.contains("b")).count();

        System.out.println("Lines with a: " + numAs + ", lines with b: " + num
        Bs);

        spark.stop();
    }
}
```

You also need to create a Maven Project Object Model (POM) file, as shown in the following example:

```
<project>
  <groupId>edu.berkeley</groupId>
  <artifactId>simple-project</artifactId>
  <modelVersion>4.0.0</modelVersion>
  <name>Simple Project</name>
  <packaging>jar</packaging>
  <version>1.0</version>
  <properties>
    <project.build.sourceEncoding>UTF-8</project.build.sourceEncoding>
    <maven.compiler.source>1.8</maven.compiler.source>
    <maven.compiler.target>1.8</maven.compiler.target>
  </properties>
  <dependencies>
    <dependency> <!-- Spark dependency -->
      <groupId>org.apache.spark</groupId>
      <artifactId>spark-sql_2.12</artifactId>
      <version>2.4.0</version>
      <scope>provided</scope>
    </dependency>
```

```
</dependencies>
</project>
```

### Before you begin

- Install Apache Spark 2.4.x, JDK 8.x, and maven
- Write a Java Spark program .java file.
- Write a pom.xml file. This is where your Scala code resides.
- If the cluster is Kerberized, ensure the required security token is authorized to compile and execute the workload.

### Procedure

1. Lay out these files according to the canonical Maven directory structure.

For example:

```
$ find .
./pom.xml
./src
./src/main
./src/main/java
./src/main/java/SimpleApp.java
```

2. Package the application using maven package command.

For example:

```
# Package a JAR containing your application
$ mvn package
...
[INFO] Building jar: {..}/{..}/target/simple-project-1.0.jar
```

After compilation, several new files are created under new directories named project and target. Among these new files, is the jar file under the target directory to run the code. For example, the file is named simple-project-1.0.jar.

3. Execute and test the workload jar using the spark submit command.

For example:

```
# Use spark-submit to run your application
spark-submit \
--class "SimpleApp" \
--master yarn \
target/simple-project-1.0.jar
```

### Compiling and running a Scala-based job

You see by example how to use sbt software to compile a Scala-based Spark job.

### About this task

In this task, you see how to use the following .sbt file that specifies the build configuration:

```
cat build.sbt
name := "Simple Project"
version := "1.0"
scalaVersion := "2.12.15"
libraryDependencies += "org.apache.spark" %% "spark-sql" % "2.4.0"
```

You also need to create a compile the following example Spark program written in Scala:

```
/* SimpleApp.scala */
import org.apache.spark.sql.Session
```

```
object SimpleApp {
  def main(args: Array[String]) {
    val logFile = "YOUR_SPARK_HOME/README.md" // Should be some file on your
system
    val spark = SparkSession.builder.appName("Simple Application").getOrCreate()
    val logData = spark.read.textFile(logFile).cache()
    val numAs = logData.filter(line => line.contains("a")).count()
    val numBs = logData.filter(line => line.contains("b")).count()
    println(s"Lines with a: $numAs, Lines with b: $numBs")
    spark.stop()
  }
}
```

### Before you begin

- Install Apache Spark 2.4.x.
- Install JDK 8.x.
- Install Scala 2.12.
- Install Sbt 0.13.17.
- Write an .sbt file for configuration specifications, similar to a C include file.
- Write a Scala-based Spark program (a .scala file).
- If the cluster is Kerberized, ensure the required security token is authorized to compile and execute the workload.

### Procedure

1. Compile the code using sbt package command from the directory where the build.sbt file exists.

For example:

```
# Your directory layout should look like this
$ find .
.
./build.sbt
./src
./src/main
./src/main/scala
./src/main/scala/SimpleApp.scala

# Package a jar containing your application
$ sbt package
...
[info] Packaging {..}/{..}/target/scala-2.12/simple-project_2.12-1.0.jar
```

Several new files are created under new directories named project and target, including the jar file named simple-project\_2.12-1.0.jar after the project name, Scala version, and code version.

2. Execute and test the workload jar using spark submit.

For example:

```
# Use spark-submit to run your application
spark-submit \
  --class "SimpleApp" \
  --master yarn \
  target/scala-2.12/simple-project_2.12-1.0.jar
```

### Running a Python-based job

You can run a Python script to execute a spark-submit or pyspark command.

**About this task**

In this task, you execute the following Python script that creates a table and runs a few queries:

```
/* spark-demo.py */
from pyspark import SparkContext
sc = SparkContext("local", "first app")
from pyspark.sql import HiveContext
hive_context = HiveContext(sc)
hive_context.sql("drop table default.sales_spark_2_copy")
hive_context.sql("CREATE TABLE IF NOT EXISTS default.sales_spark_2_copy as
select * from default.sales_spark_2")
hive_context.sql("show tables").show()
hive_context.sql("select * from default.sales_spark_2_copy limit 10").show()
hive_context.sql("select count(*) from default.sales_spark_2_copy").show()
```

**Before you begin**

Install Python 2.7 or Python 3.5 or higher.

**Procedure**

1. Log into a Spark gateway node.
2. Ensure the required security token is authorized to compile and execute the workload (if your cluster is Kerberized).
3. Execute the script using the spark-submit command.

```
spark-submit spark-demo.py --num-executors 3 --driver-memory 512m --exec
utor-memory 512m --executor-cores 1
```

4. Go to the Spark History server web UI at [http://<spark\\_history\\_server>:18088](http://<spark_history_server>:18088), and check the status and performance of the workload.

Using pyspark

**About this task**

Run your application with the pyspark or the Python interpreter.

**Before you begin**

Install PySpark using pip.

**Procedure**

1. Log into a Spark gateway node.
2. Ensure the required security token is authorized to compile and execute the workload (if your cluster is Kerberized).
3. Ensure the user has access to the workload script (python or shell script).
4. Execute the script using pyspark.

```
pyspark spark-demo.py --num-executors 3 --driver-memory 512m --executor-
memory 512m --executor-cores 1
```

5. Execute the script using the Python interpreter.

```
python spark-demo.py
```

6. Go to the Spark History server web UI at [http://<spark\\_history\\_server>:18088](http://<spark_history_server>:18088), and check the status and performance of the workload.

**Running a job interactively**

## About this task

### Procedure

1. Log into a Spark gateway node.
2. Ensure the required security token is authorized to compile and execute the workload (if your cluster is Kerberized).
3. Launch the “spark-shell”.

For example:

```
spark-shell --jars target/mylibrary-1.0-SNAPSHOT-jar-with-dependencies.jar
```

4. Create a Spark context and run workload scripts.

```
cala> import org.apache.spark.sql.hive.HiveContext
scala> val sqlContext = new HiveContext(sc)
scala> sqlContext.sql("CREATE TABLE IF NOT EXISTS default.sales_spark_1(Region string, Country string,Item_Type string,Sales_Channel string,Order_Priority string,Order_Date date,Order_ID int,Ship_Date date,Units_sold string,Unit_Price string,Unit_cost string>Total_revenue string>Total_Cost string>Total_Profit string) row format delimited fields terminated by ','")
scala> sqlContext.sql("load data local inpath '/tmp/sales.csv' into table default.sales_spark_1")
scala> sqlContext.sql("show tables")
scala> sqlContext.sql("select * from default.sales_spark_1 limit 10").show()
scala> sqlContext.sql("select count(*) from default.sales_spark_1").show()
```

5. Go to the Spark History server web UI at [http://<spark\\_history\\_server>:18088](http://<spark_history_server>:18088), and check the status and performance of the workload.

### Post-migration tasks

After the workloads are executed on Spark 2.4, validate the output, and compare the performance of the jobs with CDH/HDP cluster executions.

After the workloads are executed on Spark 2.4, validate the output, and compare the performance of the jobs with CDH/HDP cluster executions. After you perform the post migration configurations, do benchmark testing on Spark 2.4.

Troubleshoot the failed/slow performing workloads by analyzing the application event logs/driver logs and fine tune the workloads for better performance.

For more information, see the following documents:



- <https://spark.apache.org/docs/2.4.4/sql-migration-guide-upgrade.html>  
<https://spark.apache.org/releases/spark-release-2-4-0.html>  
<https://spark.apache.org/releases/spark-release-2-2-0.html>  
<https://spark.apache.org/releases/spark-release-2-3-0.html>  
<https://spark.apache.org/releases/spark-release-2-1-0.html>  
<https://spark.apache.org/releases/spark-release-2-0-0.html>

For additional information about known issues please also refer:

[Known Issues in Cloudera Manager 7.4.4 | CDP Private Cloud](#)

For additional information about known issues please also refer to:

[Known Issues in 7.4.4 |](#)

### Spark 2.3 to Spark 2.4 Refactoring

Because Spark 2.3 is not supported on , you need to refactor Spark workloads from Spark 2.3 on CDH or HDP to Spark 2.4 on .

This document helps in accelerating the migration process, provides guidance to refactor Spark workloads and lists migration. Use this document when the platform is migrated from CDH or HDP to .

### Handling prerequisites

You must perform a number of tasks before refactoring workloads.

### About this task

Assuming all workloads are in working condition, you perform this task to meet refactoring prerequisites.

### Procedure

1. Identify all the workloads in the cluster (CDH/HDP) which are running on Spark 1.6 - 2.3.
2. Classify the workloads.

Classification of workloads will help in clean-up of the unwanted workloads, plan resources and efforts for workload migration and post upgrade testing.

Example workload classifications:

- Spark Core (scala)
- Java-based Spark jobs
- SQL, Datasets, and DataFrame
- Structured Streaming
- MLlib ()
- PySpark (Python on Spark)
- Batch Jobs
- Scheduled Jobs
- Ad-Hoc Jobs
- Critical/Priority Jobs
- Huge data Processing Jobs
- Time taking jobs
- Resource Consuming Jobs etc.
- Failed Jobs

Identify configuration changes

3. Check the current Spark jobs configuration.
  - Spark 1.6 - 2.3 workload configurations which have dependencies on job properties like scheduler, old python packages, classpath jars and might not be compatible post migration.
  - In , Capacity Scheduler is the default and recommended scheduler. Follow [Fair Scheduler to Capacity Scheduler transition](#) guide to have all the required queues configured in the cluster post upgrade. If any configuration changes are required, modify the code as per the new capacity scheduler configurations.
  - For workload configurations, see the Spark History server UI [http://spark\\_history\\_server:18088/history/<application\\_number>/environment/](http://spark_history_server:18088/history/<application_number>/environment/).
4. Identify and capture workloads having data storage locations (local and HDFS) to refactor the workloads post migration.
5. Refer to [unsupported Apache Spark features](#), and plan refactoring accordingly.

### Spark 2.3 to Spark 2.4 changes

A description of the change, the type of change, and the required refactoring provide the information you need for migrating from Spark 2.3 to Spark 2.4.

#### *Empty schema not supported*

Writing a dataframe with an empty or nested empty schema using any file format, such as parquet, orc, json, text, or csv is not allowed.

Type of change: Syntactic/Spark core

Spark 1.6 - 2.3

Writing a dataframe with an empty or nested empty schema using any file format is allowed and will not throw an exception.

Spark 2.4

An exception is thrown when you attempt to write dataframes with empty schema. For example, if there are statements such as `df.write.format("parquet").mode("overwrite").save(somePath)`, the following error occurs: `org.apache.spark.sql.AnalysisException: Parquet data source does not support null data type.`

Action Required

Make sure that DataFrame is not empty. Check whether DataFrame is empty or not as follows:

```
if (!df.isEmpty) df.write.format("parquet").mode("overwrite").save("somePath")
```

#### *CSV header and schema match*

Column names of csv headers must match the schema.

Type of change: Configuration/Spark core changes

Spark 1.6 - 2.3

Column names of headers in CSV files are not checked against the against the schema of CSV data.

Spark 2.4

If columns in the CSV header and the schema have different ordering, the following exception is thrown:`java.lang.IllegalArgumentException: CSV file header does not contain the expected fields.`

Action Required

Make the schema and header order match or set `enforceSchema` to false to prevent getting an exception. For example, read a file or directory of files in CSV format into Spark DataFrame as follows: `df3 = spark.read.option("delimiter", ";").option("header", True).option("enforeSchema", False).csv(path)`

The default "header" option is true and `enforceSchema` is False.

If `enforceSchema` is set to true, the specified or inferred schema will be forcibly applied to datasource files, and headers in CSV files are ignored. If `enforceSchema` is set to false, the schema is validated against all headers in CSV files when the header option is set to true. Field names in the schema and column names in CSV headers are checked

by their positions taking into account `spark.sql.caseSensitive`. Although the default value is `true`, you should disable the `enforceSchema` option to prevent incorrect results.

### *Table properties support*

Table properties are taken into consideration while creating the table.

Type of change: Configuration/Spark Core Changes

Spark 1.6 - 2.3

Parquet and ORC Hive tables are converted to Parquet or ORC by default, but table properties are ignored. For example, the `compression` table property is ignored:

```
CREATE TABLE t(id int) STORED AS PARQUET TBLPROPERTIES (parquet.compression
'NONE')
```

This command generates Snappy Parquet files.

Spark 2.4

Table properties are supported. For example, if no compression is required, set the `TBLPROPERTIES` as follows: (`parquet.compression 'NONE'`).

This command generates uncompressed Parquet files.

Action Required

Check and set the desired `TBLPROPERTIES`.

### *Managed table location*

Creating a managed table with nonempty location is not allowed.

Type of change: Property/Spark core changes

Spark 1.6 - 2.3

You can create a managed table having a nonempty location.

Spark 2.4

Creating a managed table with nonempty location is not allowed. In Spark 2.4, an error occurs when there is a write operation, such as `df.write.mode(SaveMode.Overwrite).saveAsTable("testdb.testtable")`. The error side-effects are the cluster is terminated while the write is in progress, a temporary network issue occurs, or the job is interrupted.

Action Required

Set `spark.sql.legacy.allowCreatingManagedTableUsingNonemptyLocation` to `true` at runtime as follows:

```
spark.conf.set("spark.sql.legacy.allowCreatingManagedTableUsingNonemptyLocation", "true")
```

### *Precedence of set operations*

Set operations are executed by priority instead having equal precedence.

Type of change: Property/Spark SQL changes

Spark 1.6 - 2.3

If the order is not specified by parentheses, equal precedence is given to all set operations.

Spark 2.4

If the order is not specified by parentheses, set operations are performed from left to right with the exception that all `INTERSECT` operations are performed before any `UNION`, `EXCEPT` or `MINUS` operations.

For example, if your code includes set operations, such as `INTERSECT`, `UNION`, `EXCEPT` or `MINUS`, consider refactoring.

Action Required

Change the logic according to following rule:

If the order of set operations is not specified by parentheses, set operations are performed from left to right with the exception that all INTERSECT operations are performed before any UNION, EXCEPT or MINUS operations.

If you want the previous behavior of equal precedence then, set `spark.sql.legacy.setopsPrecedence.enabled=true`.

#### *HAVING without GROUP BY*

HAVING without GROUP BY is treated as a global aggregate.

Type of change: Property/Spark SQL changes

Spark 1.6 - 2,3

HAVING without GROUP BY is treated as WHERE. For example, `SELECT 1 FROM range(10) HAVING true` is executed as `SELECT 1 FROM range(10) WHERE true`, and returns 10 rows.

Spark 2.4

HAVING without GROUP BY is treated as a global aggregate. For example, `SELECT 1 FROM range(10) HAVING true` returns one row, instead of 10, as in the previous version.

Action Required

Check the logic where having and group by is used. To restore previous behavior, set `spark.sql.legacy.parser.havingWithoutGroupByAsWhere=true`.

#### *CSV bad record handling*

How Spark treats malformations in CSV files has changed.

Type of change: Property/Spark SQL changes

Spark 1.6 - 2.3

CSV rows are considered malformed if at least one column value in the row is malformed. The CSV parser drops malformed rows in the DROPMALFORMED mode or outputs an error in the FAILFAST mode.

Spark 2.4

A CSV row is considered malformed only when it contains malformed column values requested from CSV datasource, other values are ignored.

Action Required

To restore the Spark 1.6 behavior, set `spark.sql.csv.parser.columnPruning.enabled` to false.

#### *Spark 2.4 CSV example*

A CSV example illustrates the CSV-handling change in Spark 2.4.

In the following CSV file, the first two records describe the file. These records are not considered during processing and need to be removed from the file. The actual data to be considered for processing has three columns (jersey, name, position).

```
These are extra line1
These are extra line2
10,Messi,CF
7,Ronaldo,LW
9,Benzema,CF
```

The following schema definition for the DataFrame reader uses the option DROPMALFORMED. You see only the required data; all the description and error records are removed.

```
schema=Structtype([Structfield("jersey",StringType()),Structfield("name",StringType()),Structfield("position",StringType())])
df1=spark.read\
.option("mode","DROPMALFORMED")\
```

```
.option("delimiter",",")\
.schema(schema)\
.csv("inputfile")
df1.select("*").show()
```

Output is:

| jersey | name    | position |
|--------|---------|----------|
| 10     | Messi   | CF       |
| 7      | Ronaldo | LW       |
| 9      | Benzema | CF       |

Select two columns from the dataframe and invoke show():

```
df1.select("jersey","name").show(truncate=False)
```

| jersey                | name    |
|-----------------------|---------|
| These are extra line1 | null    |
| These are extra line2 | null    |
| 10                    | Messi   |
| 7                     | Ronaldo |
| 9                     | Benzema |

Malformed records are not dropped and pushed to the first column and the remaining columns will be replaced with null. This is due to the CSV parser column pruning which is set to true by default in Spark 2.4.

Set the following conf, and run the same code, selecting two fields.

```
spark.conf.set("spark.sql.csv.parser.columnPruning.enabled",False)
```

```
df2=spark.read\
.option("mode","DROPMALFORMED")\
.option("delimiter",",")\
.schema(schema)\
.csv("inputfile")
df2.select("jersey","name").show(truncate=False)
```

| jersey | name    |
|--------|---------|
| 10     | Messi   |
| 7      | Ronaldo |
| 9      | Benzema |

Conclusion: If working on selective columns, to handle bad records in CSV files, set `spark.sql.csv.parser.columnPruning.enabled` to false; otherwise, the error record is pushed to the first column, and all the remaining columns are treated as nulls.

### Configuring storage locations

To execute the workloads in , you must modify the references to storage locations. In , references must be changed from HDFS to a cloud object store such as S3.

### About this task

The following sample query shows a Spark 2.4 HDFS data location.

```
scala> spark.sql("CREATE TABLE IF NOT EXISTS default.sales_spark_2(Region string, Country string,Item_Type string,Sales_Channel string,Order_Priority string,Order_Date date,Order_ID int,Ship_Date date,Units_sold string,Unit_Price string,Unit_cost string>Total_revenue string>Total_Cost string>Total_Profit string) row format delimited fields terminated by ', '")
scala> spark.sql("load data local inpath '/tmp/sales.csv' into table default.sales_spark_2")
scala> spark.sql("select count(*) from default.sales_spark_2").show()
```

The following sample query shows a Spark 2.4 S3 data location.

```
scala> spark.sql("CREATE TABLE IF NOT EXISTS default.sales_spark_2(Region string, Country string,Item_Type string,Sales_Channel string,Order_Priority string,Order_Date date,Order_ID int,Ship_Date date,Units_sold string,Unit_Price string,Unit_cost string>Total_revenue string>Total_Cost string>Total_Profit string) row format delimited fields terminated by ', '")
scala> spark.sql("load data inpath 's3://<bucket>/sales.csv' into table default.sales_spark_2")
scala> spark.sql("select count(*) from default.sales_spark_2").show()
```

### Querying Hive managed tables from Spark

Hive-on-Spark is not supported on . You need to use the Hive Warehouse Connector (HWC) to query Apache Hive managed tables from Apache Spark.

To read Hive external tables from Spark, you do not need HWC. Spark uses native Spark to read external tables. For more information, see the [Hive Warehouse Connector documentation](#).

The following example shows how to query a Hive table from Spark using HWC:

```
spark-shell --jars /opt/cloudera/parcels/CDH/jars/hive-warehouse-connector-assembly-1.0.0.7.1.4.0-203.jar --conf spark.sql.hive.hiveserver2.jdbc.url=jdbc:hive2://cdhhd02.uddeпта-bandyopadhyay-s-account.cloud:10000/default --conf spark.sql.hive.hiveserver2.jdbc.url.principal=hive/cdhhd02.uddeпта-bandyopadhyay-s-account.cloud@Uddeпта-bandyopadhyay-s-Account.CLOUD
scala> val hive = com.hortonworks.hwc.HiveWarehouseSession.session(spark).build()
scala> hive.executeUpdate("UPDATE hive_acid_demo set value=25 where key=4")
scala> val result=hive.execute("select * from default.hive_acid_demo")
scala> result.show()
```

### Compiling and running Spark workloads

After modifying the workloads, compile and run (or dry run) the refactored workloads on Spark 2.4.

You can write Spark applications using Java, Scala, Python, SparkR, and others. You build jars from these scripts using one of the following compilers.

- Java (with Maven/Java IDE),
- Scala (with sbt),
- Python (pip).
- SparkR (RStudio)

### Post-migration tasks

After the workloads are executed on Spark 2.4, validate the output, and compare the performance of the jobs with CDH/HDP cluster executions.

After the workloads are executed on Spark 2.4, validate the output, and compare the performance of the jobs with CDH/HDP cluster executions. After you perform the post migration configurations, do benchmark testing on Spark 2.4.

Troubleshoot the failed/slow performing workloads by analyzing the application event logs/driver logs and fine tune the workloads for better performance.

For more information, see the following documents:

- <https://spark.apache.org/docs/2.4.4/sql-migration-guide-upgrade.html>
- <https://spark.apache.org/releases/spark-release-2-4-0.html>
- <https://spark.apache.org/releases/spark-release-2-2-0.html>
- <https://spark.apache.org/releases/spark-release-2-3-0.html>
- <https://spark.apache.org/releases/spark-release-2-1-0.html>
- <https://spark.apache.org/releases/spark-release-2-0-0.html>

For additional information about known issues please also refer:

[Known Issues in Cloudera Manager 7.4.4 | CDP Private Cloud](#)

For additional information about known issues please also refer to:

[Known Issues in 7.4.4 |](#)

## Apache Hive Expedited Migration Tasks

If you chose to expedite the Hive upgrade process by postponing migration of your tables and databases, you need to identify any problems in tables and get help with fixing those problems before migrating the tables to CDP. You then need to migrate these tables to CDP before you can use them.

### Preparing tables for migration

You download the Hive Upgrade Check tool and use it to identify problems in unmigrated tables. These problems can cause upgrade failure. It saves time to fix the problems and avoid failure. The tool provides help for fixing those problems before migrating the tables to CDP.

### About this task

You use the Hive Upgrade Check community tool to help you identify tables that have problems affecting migration. You resolve problems revealed by the Hive Upgrade Check tool to clean up the Hive Metastore before migration. If you do not want to use the Hive Upgrade Check tool, you need to perform the tasks described in the following subtopics to migrate Hive data to CDP:

- Check SERDE Definitions and Availability
- Handle Missing Table or Partition Locations
- Manage Table Location Mapping
- Make Tables SparkSQL Compatible

### Procedure

1. Obtain the Hive Upgrade Check tool.

[Download the Hive SRE Upgrade Check tool](#) from the Cloudera labs github location.

2. Follow instructions in the github readme to run the tool.

The Hive Upgrade Check (v.2.3.5.6+) will create a yaml file (hsmm\_<name>.yaml) identifying databases and tables that require attention.

### 3. Follow instructions in prompts from the Hive Upgrade Check tool to resolve problems with the tables.

At a minimum, you must run the following processes described in the github readme:

- process ID 1 Table / Partition Location Scan - Missing Directories
- process id 3 Hive 3 Upgrade Checks - Managed Non-ACID to ACID Table Migrations

#### Creating a list of tables to migrate

To run Hive Strict Managed Migration process (HSMM) after upgrading, you need to know how to create a YAML file that specifies the tables for migration.

#### Procedure

Create a YAML file in the following format:

```
---
databaseIncludeLists:
  <database name>:
    - "<table name>"
    - "<table name2>"
    - ...
  <database name 2>:
    - ...
```

For example:

```
---
databaseIncludeLists:
  tpcds_bin_partitioned_orc_10:
    - "call_center"
    - "catalog_page"
    - "catalog_returns"
    - "customer"
    - "customer_address"
  bu_raw:
    - "cc_input"
    - "geo_regions"
```

#### Migrating tables to CDP

You set a Hive property to point to your YAML list of tables you want to migrate, and then migrate the tables by manually running the Hive Strict Managed Migration process on the tables. You perform this action to use the tables in CDP.

#### About this task

In this task you set the table migration control file URL property to the path of a YAML file that lists tables and databases you want to migrate. You run the HSMM process to migrate the tables and databases using Cloudera Manager.

#### Before you begin

- You completed the upgrade to CDP.
- You created a YAML file listing databases and tables to migrate.

#### Procedure

1. In Cloudera Manager, go to Clusters Hive-on-Tez .
2. Stop the Hive-on-Tez service.
3. In Configuration , search for table migration control file URL.
4. Set the value of the Table migration control file URL property to the absolute path and file name of your YAML include list.
5. Save configuration changes.



6. Click Clusters Hive-on-Tez , and in Actions, click Migrate Hive tables for CDP upgrade.  
HSMM migrates the Hive tables listed in the YAML.
7. To prevent problems with any subsequent HSMM run, remove the value you set for Table migration control file URL, leaving the value blank.
8. Save configuration changes.
9. Start the Hive-on-Tez service.  
The YAML-specified tables and databases are migrated.

## Apache Hive Changes in CDP

You need to know where your tables are located and the property changes that the upgrade process makes. You need to perform some post-migration tasks before using Hive tables and handle semantic changes.

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.

### Hive Configuration Property Changes

You need to know the property value changes made by the upgrade process as the change might impact your work. You might need to consider reconfiguring property value defaults that the upgrade changes.

### Hive Configuration Property Values

The upgrade process changes the default values of some Hive configuration properties and adds new properties. The following list describes those changes that occur after upgrading from CDH or HDP to .

#### **datanucleus.connectionPool.maxPoolSize**

Before upgrade: 30

After upgrade: 10

#### **datanucleus.connectionPoolingType**

Before upgrade: BONECP

After upgrade: HikariCP

#### **hive.auto.convert.join.noconditionaltask.size**

Before upgrade: 20971520

After upgrade: 52428800

Exception: Preserves pre-upgrade value if old default is overridden; otherwise, uses new default.

#### **hive.auto.convert.sortmerge.join**

Before upgrade: FALSE in the old CDH; TRUE in the old HDP.

After upgrade: TRUE

#### **hive.auto.convert.sortmerge.join.to.mapjoin**

Before upgrade: FALSE

After upgrade: TRUE

#### **hive.cbo.enable**

Before upgrade: FALSE

After upgrade: TRUE

#### **hive.cbo.show.warnings**

Before upgrade: FALSE

After upgrade: TRUE

#### **hive.compactor.worker.threads**

Before upgrade: 0

After upgrade: 5

#### **hive.compute.query.using.stats**

Before upgrade: FALSE

After upgrade: TRUE

#### **hive.conf.hidden.list**

Before upgrade:

```
javax.jdo.option.ConnectionPassword,hive.server2.keystore.password,hive.metastore.dbaccess.ssl.truststore.password,fs.s3.awsAccessKeyId,fs.s3.awsSecretAccessKey,fs.s3n.awsAccessKeyId,fs.s3n.awsSecretAccessKey,fs.s3a.access.key,fs.s3a.secret.key,fs.s3a.proxy.password,dfs.adls.oauth2.credential,fs.adl.oauth2.credential,fs.azure.account.oauth2.client.secret
```

After upgrade:

```
javax.jdo.option.ConnectionPassword,hive.server2.keystore.password,hive.druid.metadata.password,hive.driver.parallel.compilation.global.limit
```

#### **hive.conf.restricted.list**

Before upgrade:

```
hive.security.authenticator.manager,hive.security.authorization.manager,hive.users.in.admin.role,hive.server2.xsrf.filter.enabled,hive.spark.client.connect.timeout,hive.spark.client.server.connect.timeout,hive.spark.client.channel.log.level,hive.spark.client.rpc.max.size,hive.spark.client.rpc.threads,hive.spark.client.secret.bits,hive.spark.client.rpc.server.address,hive.spark.client.rpc.server.port,hive.spark.client.rpc.sasl.mechanisms,hadoop.bin.path,yarn.bin.path,spark.home,bonecp.,hikaricp.,hive.driver.parallel.compilation.global.limit,_hive.local.session.path,_hive.hdfs.session.path,_hive.tmp_table_space,_hive.local.session.path,_hive.hdfs.session.path,_hive.tmp_table_space
```

After upgrade:

```
hive.security.authenticator.manager,hive.security.authorization.manager,hive.security.metastore.authorization.manager,hive.security.metastore.authenticator.manager,hive.users.in.admin.role,hive.server2.xsrf.filter.enabled,hive.security.authorization.enabled,hive.distcp.privileged.doAs,hive.server2.authentication ldap.baseDN,hive.server2.authentication.ldap.url,hive.server2.authentication.ldap.Domain,hive.server2.authentication.ldap.groupDNPattern,hive.server2.authentication.ldap.groupFilter,hive.server2.authentication.ldap.userDNPattern,hive.server2.authentication.ldap.userFilter,hive.server2.authentication.ldap.groupMembershipKey,hive.server2.authentication.ldap.userMembershipKey,hive.server2.authentication.ldap.groupClassKey,hive.server2.authentication.ldap.customLDAPQuery,hive.privilege.synchronizer.interval,hive.spark.client.connect.timeout,hive.spark.client.server.connect.timeout,hive.spark.client.channel.log.level,hive.spark.client.rpc.max.size,hive.spark.client.rpc.threads,hive.spark.client.secret.bits,hive.spark.client.rpc.server.address,hive.spark.client.rpc.server.port,hive.spark.client.rpc.sasl.mechanisms,bonecp.,hive.druid.broker.address.default,hive.druid.coordinator.address.default,hikaricp.,hadoop.bin.path,yarn.bin.path,spark.home,hive.driver.pa
```

```
parallel.compilation.global.limit,_hive.local.session.path,_hive.hdfs.session.path,_hive.tmp_table_space,_hive.local.session.path,_hive.hdfs.session.path,_hive.tmp_table_space
```

**hive.default.fileformat.managed**

Before upgrade: None

After upgrade: ORC

**hive.default.rcfile.serde**

Before upgrade: org.apache.hadoop.hive.serde2.columnar.ColumnarSerDe

After upgrade: org.apache.hadoop.hive.serde2.columnar.LazyBinaryColumnarSerDe

Not supported in Impala. Impala cannot read Hive-created RC tables.

**hive.driver.parallel.compilation**

Before upgrade: FALSE

After upgrade: TRUE

**hive.exec.dynamic.partition.mode**

Before upgrade: strict

After upgrade: nonstrict

In , accidental use of dynamic partitioning feature is not prevented by default.

**hive.exec.max.dynamic.partitions**

Before upgrade: 1000

After upgrade: 5000

In , fewer restrictions on dynamic partitioning occur than in the pre-upgrade CDH or HDP cluster.

**hive.exec.max.dynamic.partitions.pernode**

Before upgrade: 100

After upgrade: 2000

In , fewer restrictions on dynamic partitioning occur than in the pre-upgrade CDH or HDP cluster.

**hive.exec.post.hooks**

Before upgrade:

```
com.cloudera.navigator.audit.hive.HiveExecHookContext,org.apache.hadoop.hive.ql.hooks.LineageLogger
```

After upgrade: org.apache.hadoop.hive.ql.hooks.HiveProtoLoggingHook

A prime number is recommended.

**hive.exec.reducers.max**

Before upgrade: 1099

After upgrade: 1009

Exception: Preserves pre-upgrade value if old default is overridden; otherwise, uses new default

**hive.execution.engine**

Before upgrade: mr

After upgrade: tez

Tez is now the only supported execution engine, existing queries that change execution mode to Spark or MapReduce within a session, for example, fail.

**hive.fetch.task.conversion**

Before upgrade: minimal

After upgrade: more

**hive.fetch.task.conversion.threshold**

Before upgrade: 256MB

After upgrade: 1GB

Exception: Preserves pre-upgrade value if old default is overridden; otherwise, uses new default.

**hive.hashtable.key.count.adjustment**

Before upgrade: 1

After upgrade: 0.99

Exception: Preserves pre-upgrade value if old default is overridden; otherwise, uses new default.

**hive.limit.optimize.enable**

Before upgrade: FALSE

After upgrade: TRUE

**hive.limit.pushdown.memory.usage**

Before upgrade: 0.1

After upgrade: 0.04

Exception: Preserves pre-upgrade value if old default is overridden; otherwise, uses new default.

**hive.mapjoin.hybridgrace.hashtable**

Before upgrade: TRUE

After upgrade: FALSE

**hive.mapred.reduce.tasks.speculative.execution**

Before upgrade: TRUE

After upgrade: FALSE

**hive.metastore.aggregate.stats.cache.enabled**

Before upgrade: TRUE

After upgrade: FALSE

**hive.metastore.disallow.incompatible.col.type.changes**

Before upgrade: FALSE

After upgrade: TRUE

Schema evolution is more restrictive in than in CDH to avoid data corruption. The new default disallows column type changes if the old and new types are incompatible.

**hive.metastore.dml.events**

Before upgrade: FALSE

After upgrade: TRUE

**hive.metastore.event.message.factory**

Before upgrade: org.apache.hadoop.hive.metastore.messaging.json.ExtendedJSONMessageFactory

After upgrade: org.apache.hadoop.hive.metastore.messaging.json.gzip.GzipJSONMessageEncoder

**hive.metastore.uri.selection**

Before upgrade: SEQUENTIAL

After upgrade: RANDOM

#### **hive.metastore.warehouse.dir**

Before upgrade from CDH: /user/hive/warehouse

Before upgrade from HDP: /apps/hive/warehouse

After upgrade from CDH: /warehouse/tablespace/managed/hive

After upgrade from HDP: /warehouse/tablespace/managed/hive

For information about the location of old tables and new tables, which you create after the upgrade, see Changes to CDH Hive Tables or Changes to HDP Hive tables.

#### **hive.optimize.metadataonly**

Before upgrade: FALSE

After upgrade: TRUE

#### **hive.optimize.point.lookup.min**

Before upgrade: 31

After upgrade: 2

#### **hive.prewarm.numcontainers**

Before upgrade: 10

After upgrade: 3

#### **hive.script.operator.env.blacklist**

Before upgrade: hive.txn.valid.txns,hive.script.operator.env.blacklist

After upgrade: hive.txn.valid.txns,hive.txn.tables.valid.writeids,hive.txn.valid.writeids,hive.script.operator.env.blacklist

#### **hive.security.authorization.sqlstd.confwhitelist**

Before upgrade:

```
hive\auto\.*hive\.cbo\.*hive\.convert\.*hive\.exec\.dynamic\
.partition.*hive\.exec\.*\dynamic\.partitions\.*hive\.exec\.c
ompress\.*hive\.exec\.infer\.*hive\.exec\.mode.local\.*hive\.
exec\.orc\.*hive\.exec\.parallel.*hive\.explain\.*hive\.fetch.
task\.*hive\.groupby\.*hive\.hbase\.*hive\.index\.*hive\.ind
ex\.*hive\.intermediate\.*hive\.join\.*hive\.limit\.*hive\.l
og\.*hive\.mapjoin\.*hive\.merge\.*hive\.optimize\.*hive\.or
c\.*hive\.outerjoin\.*hive\.parquet\.*hive\.ppd\.*hive\.prew
arm\.*hive\.server2\.proxy\.userhive\.skewjoin\.*hive\.smbjoin
\.*hive\.stats\.*hive\.strict\.*hive\.tez\.*hive\.vectorized
\.*mapred\.map\.*mapred\.reduce\.*mapred\.output\.compression
\.codecmapped\.job\.queuenamemapped\.output\.compression\.typema
pred\.min\.split\.sizemapreduce\.job\.reduce\.slowstart\.complet
edmapsmapped\.job\.queuenamemapped\.job\.tagmapreduce\.in
put\.fileinputformat\.split\.minsizemapreduce\.map\.*mapreduce
.reduce\.*mapreduce\.output\.fileoutputformat\.compress\.codecm
apreduce\.output\.fileoutputformat\.compress\.typeoozie\.*tez\.
am\.*tez\.task\.*tez\.runtime\.*tez\.queue\.namehive\.transpo
se\.aggr\.joinhive\.exec\.reducers\.bytes\.per\.reducerhive\.cli
ent\.stats\.countershive\.exec\.default\.partition\.namehive\.ex
ec\.drop\.ignorenonexistenthive\.counters\.group\.namehive\.defa
ult\.fileformat\.managedhive\.enforce\.bucketmapjoinhive\.enforc
e\.sortmergebucketmapjoinhive\.cache\.expr\.evaluationhive\.quer
y\.result\.fileformathive\.hashtable\.loadfactorhive\.hashtable
.initialCapacityhive\.ignore\.mapjoin\.hinhive\.limit\.row\.max
\.sizehive\.mapred\.modehive\.map\.aggrhive\.compute\.query\.usi
ng\.statshive\.exec\.rowoffsethive\.variable\.substitutehive\.va
```

```

riable\.\substitute\.\depthhive\.\autogen\.\columnalias\.\prefix\.\includefuncnamehive\.\autogen\.\columnalias\.\prefix\.\labelhive\.\exec\.\check\.\crossproductshive\.\cli\.\tez\.\session\.\asynchive\.\compathive\.\exec\.\concatenate\.\check\.\indexhive\.\display\.\partition\.\cols\.\separatelyhive\.\error\.\on\.\empty\.\partitionhive\.\execution\.\enginehive\.\exec\.\copyfile\.\maxsizehive\.\exim\.\uri\.\scheme\.\whitelisthive\.\file\.\max\.\footerhive\.\insert\.\into\.\multilevel\.\dirshive\.\localize\.\resource\.\num\.\wait\.\attemptshive\.\multi\.\insert\.\move\.\tasks\.\share\.\dependencieshive\.\support\.\quoted\.\identifiershive\.\resultset\.\use\.\unique\.\column\.\nameshive\.\analyze\.\stmt\.\collect\.\partlevel\.\statshive\.\exec\.\schema\.\evolutionhive\.\server2\.\logging\.\operation\.\levelhive\.\server2\.\thrift\.\results\.\et\.\serialize\.\in\.\taskshive\.\support\.\special\.\characters\.\tablenamehive\.\exec\.\job\.\debug\.\capture\.\stacktraceshive\.\exec\.\job\.\debug\.\timeouthive\.\llap\.\io\.\enabledhive\.\llap\.\io\.\use\.\fileid\.\pathhive\.\llap\.\daemon\.\service\.\hostshive\.\llap\.\execution\.\modehive\.\llap\.\auto\.\allow\.\uberhive\.\llap\.\auto\.\enforce\.\trehive\.\llap\.\auto\.\enforce\.\vectorizedhive\.\llap\.\auto\.\enforce\.\statshive\.\llap\.\auto\.\max\.\input\.\sizehive\.\llap\.\auto\.\max\.\output\.\sizehive\.\llap\.\skip\.\compile\.\udf\.\checkhive\.\llap\.\client\.\consistent\.\splitshive\.\llap\.\enable\.\grace\.\join\.\in\.\llaphive\.\llap\.\allow\.\permanent\.\fnshive\.\exec\.\max\.\created\.\fileshive\.\exec\.\reducers\.\maxhive\.\reorder\.\nway\.\joinshive\.\output\.\file\.\extensionhive\.\exec\.\show\.\job\.\failure\.\debug\.\infohive\.\exec\.\tasklog\.\debug\.\timeouthive\.\query\.\id

```

After upgrade:

```

hive\.\auto\.\.*hive\.\cbo\.\.*hive\.\convert\.\.*hive\.\druid\.\.*hive\.\exec\.\dynamic\.\partition\.*hive\.\exec\.\max\.\dynamic\.\partitions\.*hive\.\exec\.\compress\.\.*hive\.\exec\.\infer\.\.*hive\.\exec\.\mode.local\.\.*hive\.\exec\.\orc\.\.*hive\.\exec\.\parallel\.*hive\.\exec\.\query\.\redactor\.\.*hive\.\explain\.\.*hive\.\fetch\.\task\.\.*hive\.\groupby\.\.*hive\.\hbase\.\.*hive\.\index\.\.*hive\.\index\.\.*hive\.\intermediate\.\.*hive\.\jdbc\.\.*hive\.\join\.\.*hive\.\limit\.\.*hive\.\log\.\.*hive\.\mapjoin\.\.*hive\.\merge\.\.*hive\.\optimize\.\.*hive\.\materializedview\.\.*hive\.\orc\.\.*hive\.\outerjoin\.\.*hive\.\parquet\.\.*hive\.\ppd\.\.*hive\.\prewarm\.\.*hive\.\query\.\redaction\.\.*hive\.\server2\.\thrift\.\resultset\.\default\.\fetch\.\sizehive\.\server2\.\proxy\.\userhive\.\skewjoin\.\.*hive\.\smbjoin\.\.*hive\.\stats\.\.*hive\.\strict\.\.*hive\.\tez\.\.*hive\.\vectorized\.\.*hive\.\query\.\reexecution\.\.*reexec\.\overlay\.\.*fs\.\defaultFSssl\.\client\.\truststore\.\locationdistcp\.\atomicdistcp\.\ignore\.\failuresdistcp\.\preserve\.\statusdistcp\.\preserve\.\rawattrsdistcp\.\sync\.\foldersdistcp\.\delete\.\missing\.\sourcedistcp\.\keystore\.\resourcedistcp\.\liststatus\.\threadsdistcp\.\max\.\mapsdistcp\.\copy\.\strategydistcp\.\skip\.\crcdistcp\.\copy\.\overwritdistcp\.\copy\.\appenddistcp\.\map\.\bandwidth\.\mbdistcp\.\dynamic\.\.*distcp\.\meta\.\folderdistcp\.\copy\.\listing\.\classdistcp\.\filters\.\classdistcp\.\options\.\skipcrccheckdistcp\.\options\.\mdistcp\.\options\.\numListstatusThreadsdistcp\.\options\.\mapredSslConfdistcp\.\options\.\bandwidthdistcp\.\options\.\overwritdistcp\.\options\.\strategydistcp\.\options\.\idistcp\.\options\.\p\.*distcp\.\options\.\updatedistcp\.\options\.\deletemapred\.\map\.\.*mapred\.\reduce\.\.*mapred\.\output\.\compression\.\codecmapped\.\job\.\queue\.\namemapred\.\output\.\compression\.\typemapred\.\min\.\split\.\sizemapreduce\.\job\.\reduce\.\slowstart\.\completedmapsmareduce\.\job\.\queuenamemapreduce\.\job\.\tagsmareduce\.\input\.\fileinputformat\.\split\.\minsizemapreduce\.\map\.\.*mapreduce\.\reduce\.\.*mapreduce\.\output\.\fileoutputformat\.\compress\.\codecmareduce\.\output\.\fileoutputformat\.\compress\.\typeoozie\.\.*tez\.\am\.\.*tez\.\task\.\.*tez\.\runtime\.\.*tez\.\queue\.\namehive\.\transpose\.\aggr\.\joinhive\.\exec\.\reducers\.\bytes\.\per\.\reducerhive\.\client\.\stats\.\countershive\.\exec\.\default\.\partition\.\namehive\.\exec\.\drop\.\ignoren

```

```

onexistenthive\counters\group\namehive\default\fileformat\
managedhive\enforce\bucketmapjoinhive\enforce\sortmergebucke
tmapjoinhive\cache\expr\evaluationhive\query\result\filefo
rmathive\hashtable\loadfactorhive\hashtable\initialCapacityh
ive\ignore\mapjoin\hinhive\limit\row\max\sizehive\mapred
d\modehive\map\aggrhive\compute\query\using\statshive\ex
ec\rowoffsethive\variable\substitutehive\variable\substitut
e\depthhive\autogen\columnalias\prefix\includefuncnamehive\
autogen\columnalias\prefix\labelhive\exec\check\crossprod
uctshive\cli\tez\session\asynchive\compathive\display\partit
ion\cols\separatelyhive\error\on\empty\partitionhive\exe
cution\enginehive\exec\copyfile\maxsizehive\exim\uri\sc
heme\whitelisthive\file\max\footerhive\insert\into\multile
vel\dirshive\localize\resource\num\wait\attemptshive\multi
\insert\move\tasks\share\dependencieshive\query\results
\cache\enabledhive\query\results\cache\wait\for\pending\
resultshive\support\quoted\identifiershive\resultset\use\
unique\column\namehive\analyze\stmt\collect\partlevel\st
atshive\exec\schema\evolutionhive\server2\logging\operatio
n\levelhive\server2\thrift\resultset\serialize\in\taskshi
ve\support\special\characters\tablenamehive\exec\job\debug
\capture\stacktraceshive\exec\job\debug\timeouthive\llap
\io\enabledhive\llap\io\use\fileid\pathhive\llap\daemon
\service\hostshive\llap\execution\modehive\llap\auto\all
ow\uberhive\llap\auto\enforce\treehive\llap\auto\enforce
\vectorizedhive\llap\auto\enforce\statshive\llap\auto\ma
x\input\sizehive\llap\auto\max\output\sizehive\llap\ski
p\compile\udf\checkhive\llap\client\consistent\splitshive
\llap\enable\grace\join\in\llaphive\llap\allow\permanen
t\fnshive\exec\max\created\fileshive\exec\reducers\maxhi
ve\reorder\nway\joinshive\output\file\extensionhive\exec\
.show\job\failure\debug\infohive\exec\tasklog\debug\time
outhive\query\idhive\query\tag

```

**hive.security.command.whitelist**

Before upgrade: set,reset,dfs,add,list,delete,reload,compile

After upgrade: set,reset,dfs,add,list,delete,reload,compile,llap

**hive.server2.enable.doAs**

Before upgrade: TRUE (in case of an insecure cluster only)

After upgrade: FALSE (in all cases)

Affects only insecure clusters by turning off impersonation. Permission issues are expected to arise for affected clusters.

**hive.server2.idle.session.timeout**

Before upgrade: 12 hours

After upgrade: 24 hours

Exception: Preserves pre-upgrade value if old default is overridden; otherwise, uses new default.

**hive.server2.max.start.attempts**

Before upgrade: 30

After upgrade: 5

**hive.server2.parallel.ops.in.session**

Before upgrade: TRUE

After upgrade: FALSE

A Tez limitation requires disabling this property; otherwise, queries submitted concurrently on a single JDBC connection fail or execute slower.

**hive.server2.support.dynamic.service.discovery**

Before upgrade: FALSE

After upgrade: TRUE

**hive.server2.tez.initialize.default.sessions**

Before upgrade: FALSE

After upgrade: TRUE

**hive.server2.thrift.max.worker.threads**

Before upgrade: 100

After upgrade: 500

Exception: Preserves pre-upgrade value if the old default is overridden; otherwise, uses new default.

**hive.server2.thrift.resultset.max.fetch.size**

Before upgrade: 1000

After upgrade: 10000

**hive.service.metrics.file.location**

Before upgrade: /var/log/hive/metrics-hiveserver2/metrics.log

After upgrade: /var/log/hive/metrics-hiveserver2-hiveontez/metrics.log

This location change is due to a service name change.

**hive.stats.column.autogather**

Before upgrade: FALSE

After upgrade: TRUE

**hive.stats.deserialization.factor**

Before upgrade: 1

After upgrade: 10

**hive.support.special.characters.tablename**

Before upgrade: FALSE

After upgrade: TRUE

**hive.tez.auto.reducer.parallelism**

Before upgrade: FALSE

After upgrade: TRUE

**hive.tez.bucket.pruning**

Before upgrade: FALSE

After upgrade: TRUE

**hive.tez.container.size**

Before upgrade: -1

After upgrade: 4096

**hive.tez.exec.print.summary**

Before upgrade: FALSE

After upgrade: TRUE



**hive.txn.manager**

Before upgrade: org.apache.hadoop.hive ql.lockmgr.DummyTxnManager

After upgrade: org.apache.hadoop.hive ql.lockmgr.DbTxnManager

**hive.vectorized.execution.mapjoin.minmax.enabled**

Before upgrade: FALSE

After upgrade: TRUE

**hive.vectorized.execution.mapjoin.native.fast.hashtable.enabled**

Before upgrade: FALSE

After upgrade: TRUE

**hive.vectorized.use.row.serde.deserialize**

Before upgrade: FALSE

After upgrade: TRUE

**Related Information**

[Changes to CDH Hive Tables](#)

[Changes to HDP Hive Tables](#)

**Key syntax changes**

You need to modify queries affected by changes to Hive syntax after upgrading to . Hive has changed the syntax related to `db.table` references, such as CREATE TABLE `mydb.mytable` ... . Other syntax changes involve the LOCATION clause in CREATE TABLE. Hive in supports the enhancement to CREATE TABLE that adds the MANAGEDLOCATION clause.

**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 , 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 Private Cloud Base 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));
```

## LOCATION and MANAGEDLOCATION clauses

Before upgrading, your Hive version might have supported using the LOCATION clause in queries to create either managed or external tables or databases for managed and external tables. After upgrading, Hive stores managed and external tables in separate HDFS locations. CREATE TABLE limits the use of the LOCATION clause, and consequently requires a change to your queries. Hive in also supports a new location-related clause.

### External table limitation for creating table locations

Hive assigns a default location in the warehouse for external tables—/warehouse/tablespace/external/hive. In , Hive does not allow the LOCATION clause in queries to create a managed table. Using this clause, you can specify a location only when creating external tables. For example:

```
CREATE EXTERNAL TABLE my_external_table (a string, b string)
ROW FORMAT SERDE 'com.mytables.MySerDe'
WITH SERDEPROPERTIES ( "input.regex" = "*.csv" )
LOCATION '/warehouse/tablespace/external/hive/marketing';
```

### Table MANAGEDLOCATION clause

In , Hive has been enhanced to include a MANAGEDLOCATION clause to specify the location of managed tables as shown in the following syntax:

```
CREATE (DATABASE|SCHEMA) [IF NOT EXISTS] database_name
[COMMENT database_comment]
[LOCATION external_table_path]
[MANAGEDLOCATION managed_table_directory_path]
[WITH DBPROPERTIES (property_name=property_value, ...)];
```

Hive assigns a default location in the warehouse for managed tables—/warehouse/tablespace/managed/hive. In the MANAGEDLOCATION clause, you specify a top level directory for managed tables when creating a Hive database.



**Important:** Do not use the LOCATION clause to specify the location of managed tables. This clause is only used to specify the location of external tables. Use the MANAGEDLOCATION clause if you are creating managed tables. You must also ensure that you do not set LOCATION and MANAGEDLOCATION to the same HDFS path.

Use DESCRIBE DATABASE db\_name; to view the root location of the database on the filesystem.

## Related Information

[Create a default directory for managed tables](#)

## 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

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

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 TIMESTAMP) ;
...
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

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

In , `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
- 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 , 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 .

### Before upgrading to CDP Private Cloud Base

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 Private Cloud Base

- 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:

### Modify the default CREATE TABLE behavior

#### 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 `hive-site.xml` under Hive configuration. 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 `as` 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](#).

### Related Information

[Configuring legacy CREATE TABLE behavior](#)

[Introduction to Hive Warehouse Connector](#)

[Apache Hive 3 tables](#)

[Configure a Resource-based policy: Hive](#)

[HDFS ACLs](#)

### 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 `hive-site.xml`, 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 `hive-site.xml`. When configured at the site level, the behavior persists from session to session.

## Procedure

1. In , 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.

The screenshot shows the Cloudera Manager configuration interface for the HIVE\_ON\_TEZ-1 service. The search bar at the top contains the text 'hive.create'. On the left, there are filter sections for 'SCOPE' and 'CATEGORY'. The 'SCOPE' section lists 'HIVE\_ON\_TEZ-1 (Service-Wide)' with a count of 2, 'Gateway' with 0, and 'HiveServer2' with 0. The 'CATEGORY' section lists 'Main' with 0. The main configuration area displays two properties: 'Default Table Format - Create Tables as ACID Insert Only' and 'Default Table Format - Create Tables as Full ACID'. Both properties have a checked checkbox and are associated with 'HIVE\_ON\_TEZ-1 (Service-Wide)'. The ACID property also shows the value 'hive.create.as.insert.only' and a gear icon for configuration. The Full ACID property shows the value 'hive.create.as.acid' and a gear icon.

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 table property set to true.

### Disabling Partition Type Checking

An enhancement in Hive 3 checks the types of partitions. This feature can be disabled by setting a property. For more information, see the ASF Apache Hive Language Manual.

Before Upgrade to

In CDH 5.x, partition values are not type checked.

After Upgrade to

Partition values specified in the partition specification are type checked, converted, and normalized to conform to their column types if the property hive.typecheck.on.insert is set to true (default). The values can be numbers.

Action Required

If type checking of partitions causes problems, disable the feature. To disable partition type checking, set hive.typecheck.on.insert to false. For example:

```
SET hive.typecheck.on.insert=false;
```

### 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 Private Cloud Base

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 Private Cloud Base

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

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

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

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

#### After Upgrade to

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

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 understand the interfaces that are not supported.

### Unsupported Interfaces and features

The following interfaces are not supported in :

- 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 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 . 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.

### Hive-Kudu integration

does not support the integration of HiveServer (HS2) with Kudu tables. You cannot run queries against Kudu tables from HS2.

### Partially unsupported interfaces

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

### Unsupported Features

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

- 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 .

In , 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 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

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

### Unsupported Connector Use

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](#).

### Related Information

[Configuring HMS for high availability](#)

### 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 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 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 Actions.



2. In , 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 Private Cloud Base moves managed files to the Hive warehouse / `warehouse/tablespace/managed/hive`. The upgrade process carries the external files over to CDP Private Cloud Base 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 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 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 Private Cloud Base 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 . The ownership of the Hive table file is a factor in determining table types and ACID operations after the upgrade.

**Table 5: HDP 2.x and 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      |

### Apache Hive Post-Upgrade Tasks

A successful upgrade requires performing a number of procedures that you can follow using step-by-step instructions. Important configuration tasks set up security on your cluster. You learn about semantic changes that might affect your applications, and see how to find your tables or move them. You find out about the Hive Warehouse Connector (HWC) to access files from Spark.

### Customizing critical Hive configurations

As Administrator, you need property configuration guidelines. You need to know which properties you need to reconfigure after upgrading. You must understand which the upgrade process carries over from the old cluster to the new cluster.

The upgrade process tries to preserve your Hive configuration property overrides. These overrides are the custom values you set to configure Hive in the old CDH or HDP cluster. The upgrade process does not preserve all overrides. For example, a custom value you set for `hive.exec.max.dynamic.partitions.pernode` is preserved. In the case of other properties, for example `hive.cbo.enable`, the upgrade ignores any override and just sets the -recommended value.

The upgrade process does not preserve overrides to the configuration values of the following properties that you likely need to reconfigure to meet your needs:

- `hive.conf.hidden.list`
- `hive.conf.restricted.list`
- `hive.exec.post.hooks`
- `hive.script.operator.env.blacklist`
- `hive.security.authorization.sqlstd.confwhitelist`
- `hive.security.command.whitelist`

The Apache Hive Wiki describes these properties. The values of these properties are lists.

The upgrade process ignores your old list and sets a new generic list. For example, the `hive.security.command.whitelist` value is a list of security commands you consider trustworthy and want to keep. Any overrides of this list that you set in the old cluster are not preserved. The new default is probably a shorter (more restrictive) list than the original default you were using in the old cluster. You need to customize this to meet your needs.

Check and change each property listed above after upgrading as described in the next topic.

Consider reconfiguring more property values than the six listed above. Even if you did not override the default value in the old cluster, the default might have changed in a way that impacts your work.

### Setting Hive Configuration Overrides

You need to know how to configure the critical customizations that the upgrade process does not preserve from your old Hive cluster. Referring to your records about your old configuration, you follow steps to set at least six critical property values.

#### About this task

By design, the six critical properties that you need to customize are not visible in `Visible` in column of Configurations Requirements and Recommendations. You use the Safety Valve to add these properties to `hive-site.xml` as shown in this task.

#### Procedure

1. In `Clusters` select the Hive on Tez service. Click `Configuration`, and search for `hive-site.xml`.
2. In `Hive Service Advanced Configuration Snippet (Safety Valve) for hive-site.xml`, click `+`.

The screenshot shows the Cloudera Manager interface. On the left is a navigation sidebar with options like Clusters, Hosts, Diagnostics, Audits, Charts, Replication, Administration, and Private Cloud. The main content area is titled 'Configuration' and shows a search for 'hive-site'. Below the search bar, there are filters for SCOPE and CATEGORY. The SCOPE filter shows 'HIVE\_ON\_TEZ-1 (Service-Wide)' with a count of 2, and 'Gateway' and 'HiveServer2' with counts of 1. The CATEGORY filter shows 'Advanced' with a count of 4. On the right, the configuration snippet editor is open for 'HIVE\_ON\_TEZ-1 (Service-Wide)'. It shows a table with columns for Name, Value, and Description. The Name column contains 'Hive Service Advanced Configuration Snippet (Safety Valve) for hive-site.xml'. The Value and Description columns are empty text input fields. There are also buttons for 'Undo', 'View as XML', and 'Show All Descriptions'.

3. In `Name`, add the `hive.conf.hidden.list` property.
4. In `Value`, add your custom list.

5. Customize the other critical properties: `hive.conf.restricted.list`, `hive.exec.post.hooks`, `hive.script.operator.env.blacklist`, `hive.security.authorization.sqlstd.confwhitelist`, `hive.security.command.whitelist`.  
Use `hive.security.authorization.sqlstd.confwhitelist.append`, for example, to set up the list.
6. Save the changes and restart the Hive service.
7. Look at the Configurations Requirements and Recommendations to understand which overrides were preserved or not.

### Hive Configuration Requirements and Recommendations

You need to set certain Hive and HiveServer (HS2) configuration properties after upgrading. You review recommendations for setting up for your needs, and understand which configurations remain unchanged after upgrading, which impact performance, and default values.

### Requirements and Recommendations

The following table includes the Hive service and HiveServer properties that the upgrade process changes. Other property values (not shown) are carried over unchanged from CDH or HDP to

- Set After Upgrade column: properties you need to manually configure after the upgrade to . Pre-existing customized values are not preserved after the upgrade.
- Default Recommended column: properties that the upgrade process changes to a new value that you are strongly advised to use.
- Impacts Performance column: properties changed by the upgrade process that you set to tune performance.
- Safety Value Overrides column: How the upgrade process handles Safety Valve overrides.
  - Disregards: the upgrade process removes any old CDH Safety Valve configuration snippets from the new CDP configuration.
  - Preserves means the upgrade process carries over any old CDH snippets to the new CDP configuration.
  - Not applicable means the value of the old parameter is preserved.
- Visible in CM column: property is visible in Cloudera Manager after upgrading. after upgrading.

If a property is not visible, and you want to configure it, use the Safety Valve to safely add the parameter to the correct file, for example to a cluster-wide, `hive-site.xml` file.

**Table 6:**

| Property   | Set After Upgrade | Default Recommended | Impacts Performance | New Feature | Safety Valve Overrides | Visible in CM |
|--|-------------------|---------------------|---------------------|-------------|------------------------|---------------|
| <code>datanucleus.connectionPool.maxPoolSize</code>        |                   |                     | #                   |             | Preserve               |               |
| <code>datanucleus.connectionPoolingType</code>             |                   |                     | #                   |             | Disregard              |               |
| <code>hive.async.log.enabled</code>                        |                   |                     |                     |             | Disregard              | #             |
| <code>hive.auto.convert.join.noconditionaltask.size</code> |                   |                     |                     |             | Not applicable         | #             |
| <code>hive.auto.convert.sortmerge.join</code>              |                   |                     |                     |             | Preserve               |               |
| <code>hive.auto.convert.sortmerge.join.to.mapjoin</code>   |                   |                     |                     |             | Preserve               |               |
| <code>hive.cbo.enable</code>                               |                   |                     |                     |             | Disregard              | #             |
| <code>hive.cbo.show.warnings</code>                        |                   |                     |                     |             | Disregard              |               |
| <code>hive.compactor.worker.threads</code>                 |                   |                     |                     | #           | Disregard              | #             |
| <code>hive.compute.query.using.stats</code>                |                   |                     | #                   |             | Disregard              | #             |
| <code>hive.conf.hidden.list</code>                         | #                 |                     |                     |             | Disregard              |               |
| <code>hive.conf.restricted.list</code>                     | #                 |                     |                     |             | Disregard              |               |
| <code>hive.default.fileformat.managed</code>               |                   |                     |                     |             | Disregard              | #             |
| <code>hive.default.rcfile.serde</code>                     |                   | #                   |                     |             | Preserve               |               |

| Property  | Set After Upgrade | Default Recomm          | Impacts Performa | New Feature | Safety Valve Overrides | Visible in CM |
|---|-------------------|-------------------------|------------------|-------------|------------------------|---------------|
| hive.driver.parallel.compilation                      |                   |                         |                  |             | Disregard              | #             |
| hive.exec.dynamic.partition.mode                      |                   |                         |                  |             | Disregard              |               |
| hive.exec.max.dynamic.partitions                      |                   |                         |                  |             | Preserve               |               |
| hive.exec.max.dynamic.partitions.pernode              |                   |                         |                  |             | Preserve               |               |
| hive.exec.post.hooks                                  | #                 |                         |                  |             | Disregard              |               |
| hive.exec.reducers.max                                |                   | # or other prime number |                  |             | Not applicable         | #             |
| hive.execution.engine                                 |                   |                         |                  |             | Disregard              |               |
| hive.fetch.task.conversion                            |                   |                         | #                |             | Not applicable         | #             |
| hive.fetch.task.conversion.threshold                  |                   |                         | #                |             | Not applicable         | #             |
| hive.hashtable.key.count.adjustment                   |                   |                         | #                |             | Preserve               |               |
| hive.limit.optimize.enable                            |                   | #                       |                  |             | Disregard              |               |
| hive.limit.pushdown.memory.usage                      |                   |                         | #                |             | Not Applicable         | #             |
| hive.mapjoin.hybridgrace.hashtable                    |                   | #                       | #                |             | Disregard              |               |
| hive.mapred.reduce.tasks.speculative.execution        |                   | #                       |                  |             | Disregard              |               |
| hive.metastore.aggregate.stats.cache.enabled          |                   | #                       | #                |             | Disregard              |               |
| hive.metastore.disallow.incompatible.col.type.changes |                   |                         |                  |             | Disregard              |               |
| hive.metastore.dml.events                             |                   |                         |                  |             | Disregard              | #             |
| hive.metastore.event.message.factory                  |                   | #                       |                  |             | Disregard              |               |
| hive.metastore.uri.selection                          |                   | #                       |                  |             | Disregard              |               |
| hive.metastore.warehouse.dir                          |                   |                         |                  |             | Preserve               | #             |
| hive.optimize.metadataonly                            |                   | #                       |                  |             | Disregard              |               |
| hive.optimize.point.lookup.min                        |                   |                         |                  |             | Disregard              |               |
| hive.prewarm.numcontainers                            |                   |                         |                  |             | Disregard              |               |
| hive.script.operator.env.blacklist                    | #                 |                         |                  |             | Disregard              |               |
| hive.security.authorization.sqlstd.confwhitelist      | #                 |                         |                  |             | Disregard              |               |
| hive.security.command.whitelist                       | #                 |                         |                  |             | Disregard              |               |
| hive.server2.enable.doAs                              |                   |                         |                  |             | Disregard              | #             |
| hive.server2.idle.session.timeout                     |                   |                         |                  |             | Not applicable         | #             |
| hive.server2.max.start.attempts                       |                   |                         |                  |             | Preserve               |               |
| hive.server2.parallel.ops.in.session                  |                   |                         |                  |             | Preserve               |               |
| hive.server2.support.dynamic.service.discovery        |                   |                         |                  | #           | Disregard              | #             |
| hive.server2.tez.initialize.default.sessions          |                   |                         |                  | #           | Disregard              |               |
| hive.server2.thrift.max.worker.threads                |                   |                         |                  |             | Not Applicable         | #             |
| hive.server2.thrift.resultset.max.fetch.size          |                   |                         |                  |             | Preserve               |               |
| hive.service.metrics.file.location                    |                   |                         |                  |             | Disregard              | #             |
| hive.stats.column.autogather                          |                   | #                       |                  |             | Disregard              |               |
| hive.stats.deserialization.factor                     |                   | #                       |                  |             | Disregard              |               |

| Property  | Set After Upgrade | Default Recomm | Impacts Performa | New Feature | Safety Valve Overrides | Visible in CM |
|---|-------------------|----------------|------------------|-------------|------------------------|---------------|
| hive.support.special.characters.tablename                       |                   | #              |                  |             | Disregard              |               |
| hive.tez.auto.reducer.parallelism                               |                   |                |                  | #           | Disregard              | #             |
| hive.tez.bucket.pruning   |                   |                |                  | #           | Disregard              | #             |
| hive.tez.container.size   |                   |                |                  | #           | Disregard              | #             |
| hive.tez.exec.print.summary                                     |                   |                |                  | #           | Disregard              | #             |
| hive.txn.manager  |                   |                |                  | #           | Disregard              | #             |
| hive.vectorized.execution.mapjoin.minmax.enabled                |                   | #              |                  |             | Disregard              |               |
| hive.vectorized.execution.mapjoin.native.fast.hashtable.enabled |                   | #              |                  |             | Disregard              |               |
| hive.vectorized.use.row.serde.deserialize                       |                   | #              |                  |             | Disregard              |               |

### Removing the LLAP Queue

Before upgrading from HDP to CDP, if you used LLAP, a YARN interactive query queue was created. This queue is carried over to your CDP cluster. You must remove this queue, likely named llap, after the upgrade.


#### About this task

When you set up LLAP in an HDP cluster, Ambari creates a queue named llap by default; however, you might have created the LLAP queue manually and assigned a different name to the queue. Look for your LLAP and remove it as follows:



**Note:** LLAP is not supported on Cloudera Private Cloud Base.

#### Procedure

1. In Cloudera Manager, select Clusters YARN YARN Queue Manager UI .  
A graphical queue hierarchy is displayed in the Overview tab.
2.   
Click the options menu for the interactive query queue.
3. Click Delete Queue, and confirm deletion.

### Configuring HiveServer for ETL using YARN queues

You need to set several configuration properties to allow placement of the Hive workload on the Yarn queue manager, which is common for running an ETL job. You need to set several parameters that effectively disable the reuse of containers. Each new query gets new containers routed to the appropriate queue.

#### About this task

Hive configuration properties affect mapping users and groups to YARN queues. You set these properties to use with YARN Placement Rules.

To set Hive properties for YARN queues:

#### Procedure

1. In Cloudera Manager, click Clusters Hive-on-Tez Configuration .
2. Search for the Hive Service Advanced Configuration Snippet (Safety Valve) for hive-site.xml setting.
3. In the Hive Service Advanced Configuration Snippet (Safety Valve) for hive-site.xml setting, click +.
4. In Name enter the property hive.server2.tez.initialize.default.sessions and in value enter false.
5. In Name enter the property hive.server2.tez.queue.access.check and in value enter true.
6. In Name enter the property hive.server2.tez.sessions.custom.queue.allowed and in value enter true.

### Removing Hive on Spark Configurations

Your scripts, or queries, include the Hive on Spark configuration, which is no longer supported, and you must know how to recognize and remove these configurations.

In , there is no Hive-Spark dependency. The Spark site and libs are not in the classpath. This execution engine has been replaced by Apache Tez.

Before Upgrade to

CDH supported Hive on Spark and the following configuration to enable Hive on Spark: `set hive.execution.engine=spark`

After Upgrade to

does not support Hive on Spark. Scripts that enable Hive on Spark do not work.

Action Required

Remove `set hive.execution.engine=spark` from your scripts.

### Configuring authorization to tables

Although the upgrade process makes no change to the location of external tables, you need to set up access to external tables in HDFS. If you choose the recommended Ranger security model for authorization, you need to set up policies and configure Hive metastore (HMS).

#### About this task

Set up access to external tables in HDFS using one of the following methods.

- Set up a Hive HDFS policy in Ranger (recommended) to include the paths to external table data.
- Put an HDFS ACL in place. Store the external text file, for example a comma-separated values (CSV) file, in HDFS that will serve as the data source for the external table.

If you want to use Ranger to authorize access to your tables, you must configure a few HMS properties for authorization in addition to setting up Ranger policies. If you have not configured HMS, attempting to create a table using Spark SQL, Beeline, or Hue results in the following error:

```
org.apache.hadoop.hive.ql.ddl.DDLTask. MetaException(message:No privilege 'Create' found for outputs { database:DATABASE_NAME, table:TABLE_NAME})
```

#### Related Information

[Authorizing Apache Hive Access](#)

[Configuring HMS properties for authorization](#)

#### Setting up access control lists

Several sources of information about setting up HDFS ACLs plus a brief Ranger overview and pointer to Ranger information prepare you to set up Hive authorization.

In , HDFS supports POSIX ACLs (Access Control Lists) to assign permissions to users and groups. In lieu of Ranger policies, you use HDFS ACLs to check and make any necessary changes in HDFS permission changes. For more information, see [HDFS ACLs](#), [Apache Software Foundation HDFS Permissions Guide](#), and [HDFS ACL Permissions](#).

In Ranger, you give multiple groups and users specific permissions based on your use case. You apply permissions to a directory tree instead of dealing with individual files. For more information, see [Authorizing Apache Hive Access](#).

If possible, you should use Ranger policies over HDFS ACLs to control HDFS access. Controlling HDFS access through Ranger provides a single, unified interface for understanding and managing your overall governance framework and policy design. If you need to mimic the legacy Sentry HDFS ACL Sync behavior for Hive and Impala tables, consider using Ranger RMS.

#### Configure encryption zone security

Under certain conditions, you as Administrator, need to perform a security-related task to allow users to access to tables stored in encryption zones. You find out how to prevent access problems to these tables.



### About this task

Hive on Tez cannot run some queries on tables stored in encryption zones under certain conditions. Perform the following procedure only when the cluster uses self-signed certificates.



**Important:** Skip this task for clusters where TLS certificates are properly signed by a Certificate Authority (CA), and the CA is in the truststore files.

### Procedure

1. Copy the `ssl-client.xml` file to a directory that is available on all hosts.
2. In `Clusters Hive on Tez Configuration`.
3. Search for the Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml` setting.
4. In the Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml` setting, click **+**.
5. In Name enter the property `tez.aux.uris` and in value enter `path-to-ssl-client.xml`.  
Ensure that you include the file URI scheme in the path. For example:

```
tez.aux.uris=file:///etc/hadoop/conf
```

### Configure edge nodes as gateways

If you use command-line clients, such as Sqoop, to access Hive, you must configure these gateways to use defaults for your service. You can accomplish this task in a few steps.

### About this task

By default, the HS2 instances configured in the migration already have the default `beeline-site.xml` file defined for the service. Other hosts do not. Configure these hosts as a gateway for that service.

### Procedure

1. Find the notes you made before the upgrade about edge nodes and default, connected endpoints.
2. In `Clusters Hive on Tez Configuration`, configure hosts other than HiveServer (HS2) hosts that you want to be Hive Gateway nodes as gateways for the default `beeline-site.xml` file for the gateway service.

### Spark integration with Hive

You need to know a little about Hive Warehouse Connector (HWC) and how to find more information because to access Hive from Spark, you need to use HWC implicitly or explicitly.

You can use the Hive Warehouse Connector (HWC) to access Hive managed tables from Spark. HWC is specifically designed to access managed ACID v2 Hive tables, and supports writing to tables in Parquet, ORC, Avro, or Textfile formats. HWC is a Spark library/plugin that is launched with the Spark app.

You do not need HWC to read from or write to Hive external tables. Spark uses native Spark to access external tables.

Use the Spark Direct Reader and HWC for ETL jobs. For other jobs, consider using Apache Ranger and the `HiveWarehouseConnector` library to provide row and column, fine-grained access to the data.

HWC supports `spark-submit` and `pyspark`. The spark thrift server is not supported.


### Configure HiveServer HTTP mode

If you use Knox, you might need to change the HTTP mode configuration. If you installed Knox on and want to proxy HiveServer with Knox, you need to change the default HiveServer transport mode (`hive.server2.transport.mode`).

### Procedure

1. Click `Clusters HIVE_ON_TEZ Configuration`
2. In Search, type `transport`.

- In HiveServer2 Transport Mode, select http.



- Save and restart Hive on Tez.

### Configuring HMS for high availability

To provide failover to a secondary Hive metastore if your primary instance goes down, you need to know how to add a Metastore role in Cloudera Manager and configure a property.

#### About this task

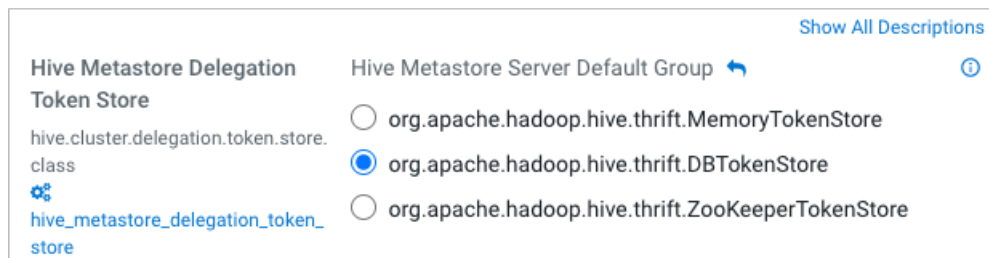
Multiple HMS instances run in active/active mode. No load balancing occurs. An HMS client always reaches the first instance unless it is down. In this case, the client scans the `hive.metastore.uris` property that lists the HMS instances for a replacement HMS. The second HMS is the designated replacement if `hive.metastore.uri.selection` is set to `SEQUENTIAL` (recommended and the default); otherwise, the replacement is selected randomly from the list if `hive.metastore.uri.selection` is set to `RANDOM`.

#### Before you begin

Minimum Required Role: Configurator (also provided by Cluster Administrator, Full Administrator)


#### Procedure

- In Cloudera Manager, click Clusters Hive Configuration .
- Take one of the following actions:
  - If you have a cluster secured by Kerberos, search for Hive Delegation Token Store, which specifies storage for the Kerberos token as described below.
  - If you have an unsecured cluster, skip the next step.
- Select `org.apache.hadoop.hive.thrift.DBTokenStore`, and save the change.



Storage for the Kerberos delegation token is defined by the `hive.cluster.delegation.token.store.class` property. The available choices are Zookeeper, the Metastore, and memory. Cloudera recommends using the database by setting the `org.apache.hadoop.hive.thrift.DBTokenStore` property.

4. Click Instances Actions Add Role Instances



5. In Assign Roles, in Metastore Server, click Select Hosts.

6. In Hosts Selected, scroll and select the host that you want to serve as the backup Metastore, and click OK.

7. Click Continue until you exit the wizard.

8. Start the Metastore role on the host from the Actions menu.

The hive.metastore.uris property is updated automatically.

9. To check or to change the hive.metastore.uri.selection property, go to Clusters Hive Configurations , and search for Hive Service Advanced Configuration Snippet (Safety Valve) for hive-site.xml.

10. Add the property and value (SEQUENTIAL or RANDOM).

### Installing Hive on Tez and adding a HiveServer role

services include Hive on Tez and Hive Metastore (HMS). Hive on Tez is a SQL query engine using Apache Tez that performs the HiveServer (HS2) role in a cluster. You must install the Hive service (HMS) before Hive on Tez and HMS in the correct order; otherwise, HiveServer will fail. You must add the HiveServer role additionally to Hive on Tez, not the Hive service; otherwise, HiveServer will fail.

#### Procedure

1. Install the HMS, designated Hive service in .

2. Install Hive on Tez

HiveServer is installed automatically during this process.

3. Accept the default, or change the Hive warehouse location for managed and external tables as described below.


#### Adding a HiveServer role

#### Procedure

1. In Cloudera Manager, click Clusters Hive on Tez .

Do not click Clusters Hive by mistake. This selects the Hive metastore and ultimately results in failure.

2. Click Actions Add Role Instances .



- Click in the HiveServer2 box to select hosts.

| <input type="checkbox"/>            | Hostname  | IP Address  | Rack     | Cores | Physical Memory | Existing Roles  |
|-------------------------------------|---|-------------|----------|-------|-----------------|---|
| <input checked="" type="checkbox"/> | nightly7x-unsecure-1.nightly7x-unsecure.root.hwx.site | 172.27.75.0 | /default | 64    | 503.6 GiB       | CCS  G  HB...  M<br>G  HS2  LB  HS<br>AP  ES  HM  RM<br>SCM  QS  SRS  SS<br>G  G  JHS  RM |
| <input type="checkbox"/>            | nightly7x-unsecure-2.nightly7x-                       | 172.27.75.2 | /default | 64    | 503.6 GiB       | RS  DN  G  G<br>NM  SS  G  G  |

- In the Host name column, select a host for the HiveServer2 role, and click OK. The selected host name you assigned the HiveServer2 role appears under HiveServer2.

### Assign Roles

You can specify the role assignments for your new roles here.

You can also view the role assignments by host. [View By Host](#)

Gateway × 4 HiveServer2 × (1 + 1 New)

Select hosts nightly7x-unsecure-2.nightly7x-unsecur...

- Click Continue. The new HiveServer2 role state is stopped.
- Select the new HiveServer2 role.

| Actions for Selected (1) ▾          |        |             |         |                       |
|-------------------------------------|--------|-------------|---------|-----------------------|
| <input type="checkbox"/>            | Status | Role Type   | State   | Hostname              |
| <input checked="" type="checkbox"/> |        | HiveServer2 | Stopped | nightly7x-unsecure-2. |
| <input type="checkbox"/>            |        | HiveServer2 | Started | nightly7x-unsecure-1. |

- In Actions for Selected, select Start, and then click Start to confirm. You see that the service successfully started.

### Start

Status **Finished** Context [HiveServer2 \(nightly7x-unsecure-2\)](#) Apr 1, 3:08:53 AM

23.15s

Successfully started service.

✓ **Completed 1 of 1 step(s).**

Show All Steps
  Show Only Failed Steps
  Show Only Running Steps

> Starting 1 roles on service

### Changing the Hive warehouse location


#### About this task

You use the Hive Metastore Action menu in Cloudera Manager, and 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.

## Updating Hive and Impala JDBC/ODBC drivers

After upgrading, Cloudera recommends that you update your Hive and Impala JDBC and ODBC drivers. You follow a procedure to download a driver.

### Before you begin

Configure authenticated users for running SQL queries through a JDBC or ODBC driver. For example, set up a Ranger policy.

### Getting the JDBC driver

You learn how to download the Cloudera Hive and Impala JDBC drivers to give clients outside the cluster access to your SQL engines.

## Procedure

1. Download the latest Hive JDBC driver for CDP from the [Hive JDBC driver download page](#).
2. Go to the [Impala JDBC driver](#) page, and download the latest Impala JDBC driver.
3. Follow JDBC driver installation instructions on the download page.

### Getting the ODBC driver

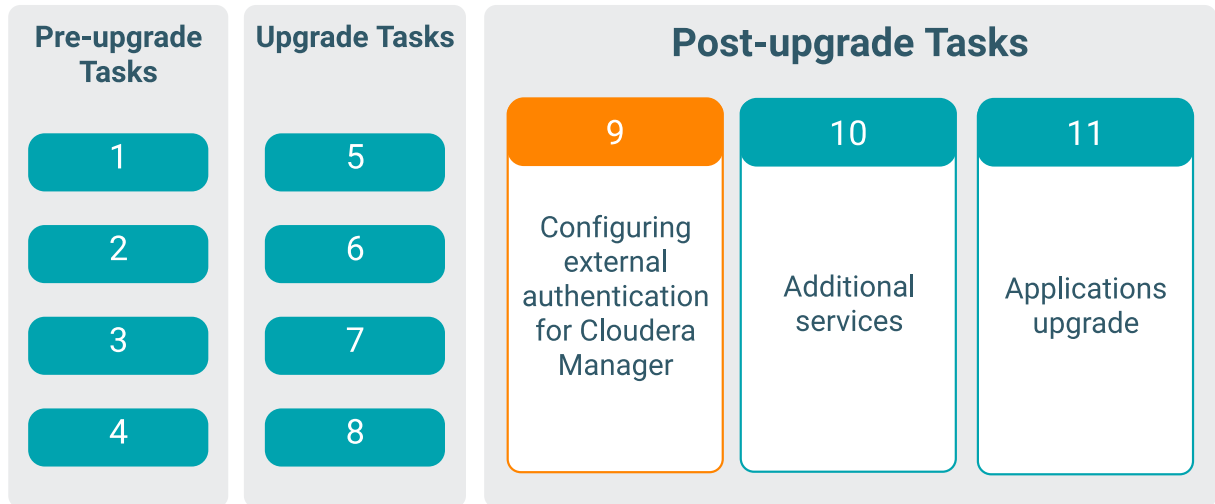
You learn how to download the Cloudera ODBC drivers for Hive and Impala.

## Procedure

1. Download the latest Hive ODBC driver for CDP from the [Cloudera ODBC driver download page](#).
2. Go to the [Impala ODBC driver](#) page, and download the latest Impala ODBC driver.
3. Follow ODBC driver installation instructions on the download page.

# Configuring External Authentication for Cloudera Manager


If you have LDAP configured on your cluster, then continue further with this document.



[Configuring External Authentication](#)

## Additional Services

The Hue and DAS services are optional.



## Installing Data Analytics Studio on a Cloudera Private Cloud Base cluster using Ambari

This section applies only if you were on HDP 2.6.5 running Hive LLAP and Druid workloads, and using Ambari Views to manage and monitor the performance of your queries. With Cloudera Private Cloud Base, Data Analytics Studio (DAS) replaces Ambari Views (Hive View and Tez View) and offers additional analytical capabilities.

The tasks listed in this section assume that you have already upgraded Ambari and HDP cluster from 2.6.5 to Cloudera Private Cloud Base 7.1.x.



### Note:

1. DAS requires a backend database for storing query event information, and supports only PostgreSQL versions 9.6, 10, and 12.
2. Data Analytics Studio (DAS) has been deprecated and is no longer available in CDP Private Cloud Base starting with 7.1.9. For more information, see [Deprecation Notice for DAS](#).

## Check cluster configuration for Hive and Tez

Before you add the DAS service to your Cloudera Private Cloud Base cluster, check the configuration settings for Hive and Tez in the Ambari UI. This can ensure that your queries appear properly when you use DAS.

### Procedure

1. Go to Ambari Services Hive CONFIGS ADVANCED . Make sure that the Hive configurations are as follows:

- `hive.hook.proto.base-directory: {hive_metastore_warehouse_external_dir}/sys.db/query_data/`
- `hive.exec.failure.hooks: org.apache.hadoop.hive.ql.hooks.HiveProtoLoggingHook`
- `hive.exec.post.hooks: org.apache.hadoop.hive.ql.hooks.HiveProtoLoggingHook`
- `hive.exec.pre.hooks: org.apache.hadoop.hive.ql.hooks.HiveProtoLoggingHook`
- `hive.metastore.transactional.event.listeners: org.apache.hive.hcatalog.listener.DbNotificationListener`



**Note:** If you have upgraded your cluster from HDP 2.6.x to Cloudera Private Cloud Base 7.1.x, then specify `org.apache.hive.hcatalog.listener.DbNotificationListener` in the `hive.metastore.transactional.event.listeners` property under Advanced `hivemetastore-site`.

2. Go to Ambari Services Tez CONFIGS Custom tez-site . Make sure that the Tez configuration is as follows:

- `tez.history.logging.service.class: org.apache.tez.dag.history.logging.proto.ProtoHistoryLoggingService`
- `tez.history.logging.proto-base-dir: {hive_metastore_warehouse_external_dir}/sys.db/`

## Add the DAS service

Adding the DAS service from the Ambari UI makes DAS available on your cluster.

### Procedure

1. On the Ambari UI, click **Actions Add Service** .  
The Add Service Wizard is displayed.
2. On the Choose Services page of the Wizard, select **Data Analytics Studio**.  
Other required services are automatically selected.
3. When prompted to confirm the addition of dependent services, give a positive confirmation to all.  
This adds the other required services.

4. On the Customize Services page, expand `Advance_data_analytics_studio-database` and fill in the database details and other required fields that are highlighted.

**Note:**

- Enter a password in the password field. Do not leave this field empty.
  - Uncheck the Create Data Analytics Studio database checkbox, and specify the same database hostname in the Data Analytics Studio database hostname field on which your database is running.
5. In the `user_authentication` field, specify the authentication method that you want to use. You can specify one of the following:
    - `NONE`: Specify `NONE` if you do not want to enable user authentication. In this case, all the queries will be executed as the hive service user.
    - `KNOX_SSO`: Specify `KNOX_SSO` to enable single sign-on experience to authenticate users as configured in the Knox SSO topology.
    - `KNOX_PROXY`: Specify `KNOX_PROXY` to use DAS behind a Knox proxy. The user is authenticated as configured in the Knox proxy topology. It also enables SPNEGO authentication.
    - `SPNEGO`: Specify `SPNEGO` to authenticate users using a token-based authentication.

You can configure the user authentication separately, after you have added the DAS service as part of the post-installation tasks.

6. You can add additional configuration information in rest of the fields as appropriate.
7. Complete the remaining installation wizard steps and exit the wizard.
8. Ensure that all components required for the DAS service start successfully.
9. Make sure to restart all the affected services in Ambari.

## DAS post-installation tasks

After adding the DAS service on the Cloudera Private Cloud Base 7.1.x cluster, complete the tasks in this section to start using DAS.

### Additional configuration tasks

You can increase or decrease the time for which the audit and event logs are retained in the system, after which they are auto-purged. You can also make DAS work with HiveServer2 in case your Hive deployment is not LLAP-enabled and does not have Hive Server Interactive.

### Making DAS work with HiveServer2

If your Hive deployment is not LLAP-enabled and does not have Hive Server Interactive, you need to update the configuration for DAS to work with HiveServer2 as follows:

Go to `Ambari Data Analytics Studio Configs` and change `use.hive.interactive.mode=true` to `use.hive.interactive.mode=false`

### Customizing the retention period of audit and event logs

Audit and event logs are written by Hive/Tez to HDFS for DAS to consume to generate DAG information. Once DAS consumes them, this information on HDFS is no longer required. They are cleaned automatically via a thread in HMS (Hive metastore).

The time to live count of the audit logs is defined in the `hive.hook.proto.events.ttl` parameter. By default, the logs are retained for seven days. To configure a custom value:

1. From the Ambari UI, go to `Hive HiveMetaStore Configs` and specify the time to live in the `hive.hook.proto.events.ttl` parameter.

For example, to retain the logs for 30 days, specify `30d`.

Typically, a 30-day duration is useful if you want to download older DAG data.

2. Restart the Hive service.



### Setting up the tmp directory

DAS allows you to download logs for secure clusters. To download logs, make sure that the DAS service user has write permission to the /tmp directory. Also make sure that the /tmp directory has sufficient storage space to hold logs from a query for the download logs feature to work.

### Configuring DAS for SSL/TLS

If your HDP cluster is SSL enabled, then you can configure SSL. This is an optional configuration.

You can use one of the two options to set up SSL certificates:

- Set up trusted CA certificates
- Set up self-signed certificates

#### Set up trusted CA certificate

You can enable SSL for the DAS Engine using a certificate from a trusted Certificate Authority (CA). Certificates from a trusted CA are primarily used in production environments. For a test environment, you can use a self-signed certificate.

#### Before you begin

- You must have root user access to the clusters on which DAS Engine is installed.
- You must have obtained a certificate from your CA, following their instructions.

#### Procedure

1. Log in as root user on the cluster with DAS Engine installed.
2. Import the Certificate Chain Certificate and the certificate you obtained from your CA.

```
keytool -import -alias root -keystore <path_to_keystore_file> -trustcacerts -file <certificate_chain_certificate>
```

```
keytool -import -alias jetty -keystore <path_to_keystore_file> -file <certificate_from_CA>
```



**Note:** Ignore the following warning:

```
The JKS keystore uses a proprietary format. It is recommended to migrate to PKCS12 which is an industry standard format using "keytool -importkeystore -srckeystore <keystore_file_path> -destkeystore <keystore_file_path> -deststoretype pkcs12".
```

#### Set up self-signed certificates

You can enable SSL for the DAS Engine using a self-signed certificate. Self-signed certificates are primarily used in test environments. For a production environment, you should use a certificate from a trusted CA.

#### Before you begin

You must have root user access to the clusters on which DAS Engine is installed.

#### Procedure

1. Log in as root user on the cluster with DAS Engine installed.

## 2. Generate a key pair and keystore for use with DAS Engine.

```
keytool -genkey -alias jetty -keystore <certificate_file_path> -storepass
<keystore_password> -dname 'CN=das.host.com, OU=Eng, O=ABC Corp, L=Santa
Clara, ST=CA, C=US' -keypass <key_password> -keyalg RSA
```



**Note:** Ignore the following warning:

```
The JKS keystore uses a proprietary format. It is recommended
to migrate to PKCS12 which is an industry standard format using
"keytool -importkeystore -srckeystore <keystore_file_path> -
destkeystore <keystore_file_path> -deststoretype pkcs12".
```

Follow the prompts and enter the required information.

- CN must be the FQDN of the DAS Engine host.
- Default value for the key password is *password*.

If you change the password, then you have to update the DAS configuration.

Following is a sample command output:

```
keytool -genkey -alias jetty -keystore ~/tmp/ks -storepass password
What is your first and last name?
[Unknown]: das.host.com
What is the name of your organizational unit?
[Unknown]: Eng
What is the name of your organization?
[Unknown]: ABC Corp
What is the name of your City or Locality?
[Unknown]: Santa Clara
What is the name of your State or Province?
[Unknown]: CA
What is the two-letter country code for this unit?
[Unknown]: US
Is CN=das.host.com, OU=Eng, O=ABC Corp, L=Santa Clara, ST=CA, C=US correct?
[no]: yes

Enter key password for <jetty>
(RETURN if same as keystore password):
```



**Note:** You will have to use this keystore file while configuring the DAS Engine for TLS in Ambari.

### Configure SSL/TLS in Ambari

In the Ambari UI, you enable TLS for DAS Engine and update the DAS Engine configuration if settings change.

#### Procedure

1. Copy the keystore files generated in the earlier procedures to webapp and event processor hosts. Make sure they are owned by configured user for DAS. The default user is hive.

For example:

```
/etc/security/certs/das-cert.jks
```

2. Navigate to Data Analytics Studio Configs .
3. Set the following properties in Advanced data\_analytics\_studio-security-site section.

| Field       | Value                    |
|-------------|--------------------------|
| ssl_enabled | Make sure it is checked. |

| Field                                 | Value  |
|---------------------------------------|--|
| webapp_keystore_file                  | Enter the keystore path on the webapp host.        |
| das_webapp_keystore_password          | Enter the password used in the previous procedure. |
| event_processor_keystore_file         | Enter the keystore path on the event processor.    |
| das_event_processor_keystore_password | Enter the password used in the previous procedure. |

- In the Advanced data\_analytics\_studio-webapp-properties section, set Data Analytics Studio Webapp server protocol property to https.
- In the Advanced data\_analytics\_studio-event\_processor-properties section, set Data Analytics Studio Event Processor server protocol property to https.

### Configuring user authentication in Ambari

You can authenticate the users by using either Knox SSO, Knox proxy, or SPNEGO. You can also choose not to set up user authentication. In this case, all the queries are executed as a Hive service user.

### Configuring user authentication using Knox SSO

In a Cloudera Private Cloud Base deployment, to enable DAS to work with the CDP cluster SSO, configure the Knox settings as described in this topic.

### About this task



**Note:** Follow these instructions only if you choose to configure secure clusters.

### Before you begin

You need to export the Knox certificate from the Knox gateway host. To find the Knox gateway host, go to Ambari Services Knox CONFIGS Knox Gateway hosts .

### Procedure

- SSH in to the Knox gateway host with a root or a knoxuser user.
- Export the Knox certificate by running the following command:

```
/usr/hdp/current/knox-server/bin/knoxcli.sh export-cert --type PEM
```



**Note:** If you have already integrated Knox SSO earlier, then the gateway-identity.pem file would exist. Check whether the gateway-identity.pem file exists or not before running this command.

```
/usr/hdp/current/knox-server/data/security/keystores/gateway-identity.pem
```

If the export is successfully, the following message is displayed:

```
Certificate gateway-identity has been successfully exported to: /usr/hdp/current/knox-server/data/security/keystores/gateway-identity.pem
```

Note the location where you save the gateway-identity.pem file.

- If not done already, specify KNOX\_SSO in the user\_authentication field under Ambari Data Analytics Studio Config .



```
<rule dir="OUT" name="DAS/das/outbound/assets" pattern="assets/{**}">
  <rewrite template="{ $frontend[path] }/das/assets/{**}"/>
</rule>
<rule dir="OUT" name="DAS/das/outbound/rooturl">
  <rewrite template="{ $frontend[path] }/das/">
</rule>
<filter name="DAS/das/outbound/dasrooturl">
  <content type="application/javascript">
    <apply path="/dasroot/.." rule="DAS/das/outbound/rooturl"/>
  </content>
</filter>
</rules>
```

service.xml

```
<!--
Licensed to the Apache Software Foundation (ASF) under one or more
contributor license agreements. See the NOTICE file distributed with
this work for additional information regarding copyright ownership.
The ASF licenses this file to You under the Apache License, Version 2.0
(the "License"); you may not use this file except in compliance with
the License. You may obtain a copy of the License at

    http://www.apache.org/licenses/LICENSE-2.0
Unless required by applicable law or agreed to in writing, software
distributed under the License is distributed on an "AS IS" BASIS,
WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or impl
ied.
See the License for the specific language governing permissions and
limitations under the License.
-->
<service role="DAS" name="das" version="1.4.5">
  <routes>
    <route path="/das" />
    <route path="/das/**" />
    <route path="/das/**/hivestudio*.js">
      <rewrite apply="DAS/das/outbound/dasrooturl" to="response.body"/>
    </route>
  </routes>
</service>
```

If you have more than one Knox servers in a high availability configuration, then you need to create these files on every machine.

4. Restart the Knox server.
5. Sign in to the Knox Admin UI. From Ambari, go to Knox Quick Links and click Knox Admin UI.
6. Create or use an existing Provider Configuration with the following configuration:
  - a) (Required) Authentication provider: It provides authentication mechanism for Knox. For example, ShiroProvider. ShiroProvider integrates with LDAP.
  - b) (Required) Identity Assertion provider: It authenticates the user name requesting access to DAS through SPNEGO.
  - c) (Optional) Add any other provider as per your organization's requirements, such as Acls or WebAppSec.
7. Create a new descriptor for the DAS service with the Provider Configuration that you created earlier. If a descriptor is already present with the required Provider Configuration, then you can edit it.
  - a) Add a new service named DAS.
  - b) Under Descriptor Detail Services DAS, add the DAS Webapp URL in the URLs field and save the changes.
8. If not done already, specify KNOX\_PROXY in the user\_authentication field under Ambari Data Analytics Studio Config.

9. Next, go to Ambari Data Analytics Studio CONFIGS Advanced data\_analytics\_studio-security-site and configure the following fields:
  - a) `das_knox_user`: It specifies the Knox service user, and is picked up automatically.
  - b) `das_knox_doas_param_name`: Specify the parameter name that you specified under the Params field in the Knox Admin UI. However, if you have selected the default Identity Assertion provider, then the `das_knox_doas_param_name` name is `doAs`, and is configured automatically.
  - c) `knox_ssout_url`: Specify the KnoxSSOUT URL in the following format:  
`https://{host}:{port}/{cluster-name}/{knox-sso-out-topo}/api/v1/webssout`  
 This step requires you to have a KnoxSSO out topology present in your Knox server. To create the Knox SSO out topology, see [Enabling logout option for secure clusters](#).
  - d) Add the list of users in the `admin_users` field who need admin access to DAS.  
 You can specify \* (asterisk) in the `admin_users` field to make all users the admin users.  
 You can also specify an admin group in the `admin_groups` field.



**Note:** Only admin users have access to all the queries. Non-admin users can access only their queries.

10. Restart the DAS service.

### Configuring user authentication using SPNEGO

SPNEGO uses a Kerberized environment for user authentication. To use SPNEGO to authenticate users, specify SPNEGO in the `user_authentication` field.

#### Procedure

1. From the Ambari UI, go to Data Analytics Studio CONFIGS Advanced data\_analytics\_studio-security-site .
2. Specify SPNEGO in the `user_authentication` field.
3. Configure the `das_spnego_name_rules` field if you have specific name mapping rules for your cluster.
4. Save the changes and restart DAS.

### Enabling logout option for secure clusters

To sign out from the application (DAS Webapp) and the identity provider on secure clusters, you need to create a KnoxSSO out topology in the Ambari console and configure the `knox_ssout_url` through the Ambari UI.

#### About this task

This section can be used to configure the Knox SSO logout URL for both HA and non-HA clusters.

#### Procedure

1. SSH in to the Ambari console on which you have configured Knox as a root user.
2. If you are configuring KnoxSSOUT for the first time, then you need to create a KnoxSSO out topology called `knoxssout.xml` in the `/etc/knox/x.x.x.x-xx/0/topologies` directory.

If you have already configured KnoxSSOUT, then you can skip this step.

Add the following lines in the `knoxssout.xml` file:

```
<?xml version="1.0"?>
<topology>
  <gateway>
    <provider>
      <role>hostmap</role>
      <name>static</name>
      <enabled>>true</enabled>
    </provider>
  </gateway>
  <service>
    <role>KNOXSSOUT</role>
```

```
</service>
</topology>
```



**Note:** If you have configured Knox SSO for HA clusters, then you must create the KnoxSSO out topology on all the Knox hosts.

3. On the Ambari UI, go to `Data Analytics Studio Config Advanced data_analytics_studio-security-site` and specify the Knox SSO logout URL in the `knox_ssout_url` field in the following format:

```
https://{host}:{port}/{cluster-name}/{knox-ssout-topo}/api/v1/webssout
```

Where, `knox-ssout-topo` is the name of the KnoxSSO out topology xml file that you created earlier. In this case, it is `knoxssout`.

The host and the port that you specify here should be the same as those specified in the `knox_ssout_url` field.

To obtain the `cluster_name`, on the Ambari UI, click `admin Manage Ambari Cluster Information` and note the value from the `Cluster Name` field.

If you have HA clusters, then specify the host and the port of the load balancer pointing to the Knox instance.

4. Click `Save` and click through the confirmation pop-ups.
5. Restart DAS and any services that require restart by clicking `Actions Restart All Required`.

## Troubleshooting DAS installation

This section provides you a list of possible causes and solutions to debug and resolve issues that you might face while using DAS on a Cloudera Private Cloud Base cluster managed through Ambari.

### Problem area: Queries page

Issues related to queries can arise mostly because of misconfigured Hive or Tez hooks. You may also experience trouble viewing your queries if you are a non-admin user. Non-admin users can view only their queries.

#### Your queries are not appearing on the Queries page

Check the role that you have been assigned. If you are not an administrator, then you can only view the queries that you have run. Only administrators can view queries and data of other users.

- To enlist yourself as the admin user, go to `Ambari Data Analytics Studio Config Advanced data_analytics_studio-security-site`, and add your user name in the `admin_users` field. For example, `hive,admin1,admin2`.

You can also specify `*` (asterisk) in the `admin_users` field to make all users administrators.

You can also specify an admin group in the `admin_groups` field.

- Check the configuration settings for Hive and Tez in the Ambari UI. See [Check cluster configuration for Hive and Tez](#) on page 303.

- Verify whether the logs are added to the paths specified in the `hive.hook.proto.base-directory` and the `tez.history.logging.proto-base-dir` directory configurations.

Use the following commands to get the paths to the logs for Hive and Tez. If you have not run any queries, then the log directories may be empty.

For Hive:

```
hdfs dfs -ls /warehouse/tablespace/external/hive/sys.db/query_data
```

For Tez:

```
hdfs dfs -ls /warehouse/tablespace/external/hive/sys.db/dag_data
```

```
hdfs dfs -ls /warehouse/tablespace/external/hive/sys.db/dag_meta
```

```
hdfs dfs -ls /warehouse/tablespace/external/hive/sys.db/app_data
```

- Verify whether the following proto directory configurations are present in the `/etc/das/conf/das-event-processor.json` file: `hive.hook.proto.base-directory` and `tez.history.logging.proto-base-dir`. If these configurations are not present in the `/etc/das/conf/das-event-processor.json` file, then do not edit the `das-event-processor.json` file - because they will be overwritten by the configurations specified in Ambari.

At this stage, restart the DAS service and recheck. If this doesn't fix the problem, then you might be using multiple configuration groups in Hive and Tez.

If you are using different or multiple configuration groups in Hive or Tez, then make sure that you specify the same paths in all the configuration groups. Presently, DAS does not support multiple configuration groups.

- Verify whether you have the required access and permission to directories specified in the `hive.hook.proto.base-directory` and `tez.history.logging.proto-base-dir` directory configurations by checking errors from `EventProcessorPipeline` in the `/var/log/das/event-processor.log` file.
- Make sure that the DAS services are running as the hive service user.
- If you have a large number of users running Hive queries simultaneously using DAS, the connections from Hive can get exhausted. This sends the query to a queue in Hive which takes a long time to start running. This can result in the following:
  - If query is fired from DAS, it does not return any log or indicate that the query has been queued.
  - The query does not appear on the UI, because no event is received by DAS.

Evaluate your cluster usage. You may need to add more machines to the clusters to resolve this issue.

### Query column is empty, yet you can see the DAG ID and Application ID

- Check the configuration settings for Hive and Tez in the Ambari UI. See [Check cluster configuration for Hive and Tez](#) on page 303.
- Check whether Hive has access to write in to the directory specified by the following proto directory configuration: `hive.hook.proto.base-directory`.
- Check whether Hive can write in to the directory specified by the following proto directory configuration: `hive.hook.proto.base-directory`.
- Check whether Hive has access to read from the directory specified by the following proto directory configuration: `hive.hook.proto.base-directory`.

### Query column is not empty, but you cannot see the DAG ID and Application ID

- Check the configuration settings for Hive and Tez in the Ambari UI. See [Check cluster configuration for Hive and Tez](#) on page 303.
- Check whether the Tez job has access to write in to the directory specified by the following proto directory configuration: `tez.history.logging.proto-base-dir`.
- Check whether the data is being written in to the following subdirectories of the directory specified in the `tez.history.logging.proto-base-dir` directory configuration: `app_data`, `dag_data`, and `dag_meta`.
- Check whether das has access to read from the following subdirectories of the directory specified in the `tez.history.logging.proto-base-dir` directory configuration: `app_data`, `dag_data`, and `dag_meta`.



### You cannot view queries from other users

If you need to view queries and data of other users, then assign admin role to the user by going to Ambari Data Analytics Studio Config Advanced data\_analytics\_studio-security-site , and specifying the user in the admin\_users field. For example, hive,admin1,admin2.

You can also specify \* (asterisk) in the admin\_users field to make all users administrators.

You can also specify an admin group in the admin\_groups field.

### Problem area: Compose page

Misconfigured repl in Hive can affect the query compose page on the DAS UI. Issues can also occur if HSI or HS2 is not running.

#### You cannot see your databases or the query editor is missing

- Verify whether HSI or HS2 services are running by going to Ambari Hive Summary .
- If you are not running llap, then go to Ambari Data Analytics Studio Config Advanced data\_analytics\_studio-properties Data Analytics Studio config file template and set the use.hive.interactive.mode parameter to false.
- Verify whether the Hive configuration parameter hive.metastore.transactional.event.listeners is set to org.apache.hive.hcatalog.listener.DbNotificationListener by going to Ambari Hive Configs Advanced Advanced hive-site .
- DAS event processor runs the `repl dump` command as a hive service user. The hive service user can run the `repl dump` command if the user has repl admin privileges.
- Verify whether DAS can read the repl root directory permissions - if das and hive are running as different service users.

There may be exceptions in the `/var/log/das/event-processor.log` file with repl dump failures because of read errors.



**Note:** If the das services cannot read the repl root directory permissions, then it can cause a number of repl directories to be created, but not deleted.

#### You cannot view new databases and tables, or cannot see changes to existing databases or tables

- Verify whether the hive configuration parameter hive.metastore.transactional.event.listeners is set to org.apache.hive.hcatalog.listener.DbNotificationListener by going to Ambari Hive Configs Advanced Advanced hive-site .
- Verify whether HSI or HS2 services are running by going to Ambari Hive Summary .
- This can also happen if the Hive metadata replication fails. To troubleshoot replication-related failures, see [Replication failure in the DAS Event Processor](#) on page 313.

#### Replication failure in the DAS Event Processor

DAS uses replication as a way to copy database and table metadata information from Hive to DAS Postgres database. If the replication fails, then you may not be able to see database or table information in DAS.

The Hive metadata replication process occurs in the following two phases:

- Bootstrap dump
- Incremental dump

When the DAS Event Processor is started for the first time, the entire Hive database and table metadata is copied in to DAS. This is known as the bootstrap dump. After this phase, only the differences are copied in to DAS at one-minute intervals, from the time the last successful dump was run. This is known as an incremental dump.

If the bootstrap dump never succeeded, then you may not see any database or table information in DAS. If the bootstrap dump fails, then the information regarding the failure is captured in the most recent Event Processor log.

If an incremental dump fails, then you may not see any new changes to the databases and tables in DAS. The incremental dump relies on events stored in the Hive metastore, because these events take up a lot of space and are only used for replicating data. The events are removed from Hive metastore daily, which can affect DAS.

## Fixing incremental dump failures



**Note:** You must be an admin user to complete this task.

If you see the message “Notification events are missing in the meta store”, then reset the Postgres database using the following command:

```
curl -H 'X-Requested-By: das' -H 'Cookie: JSESSIONID=<session id cookie>' http(s)://<hostname>:<port>/api/replicationDump/reset
```

Where:

- *SESSION ID COOKIE* is the cookie value which you have to get for an admin user on the DAS UI, from the browser
- *HOSTNAME* is the DAS Webapp hostname
- *PORT* is the DAS Webapp port



**Note:**

The error message "Notification events are missing in the meta store" is a symptom and not the cause of an issue, at most times. You can usually fix this by resetting the database. However, to determine the actual cause of the failure, we need to look at the first repl dump failure. Older Event Processor logs may contain the information about the actual cause, but they are purged from the system at a set frequency. Therefore, if you hit this issue, we recommend that you first copy all the Event Processor logs that are available; else it might be too late to diagnose the actual issue.

If the first exception that you hit is an SQLException, then it is a Hive-side failure. Save the HiveServer and the Hive Metastore logs for the time when the exception occurred.

File a bug with Cloudera Support along with the above-mentioned logs.

## Problem area: Reports page

Only the admin users can view the **Reports** page. If you are a non-admin user, then you lack the privileges to view the **Reports** page.

## DAS service installation fails with the “python files missing” message

The installation might fail while you are adding the DAS service through the Ambari UI with an error such as the following: Python files are missing.

### Procedure

1. SSH in to the Ambari host as a root user.
2. Set the agent.auto.cache.update property to true in the /etc/ambari-server/conf/ambari.properties file.
3. Restart ambari using the following command:

```
ambari-server restart
```

## DAS does not log me out as expected, or I stay logged in longer than the time specified in the Ambari configuration

The login session timeout factor in DAS is governed by the following two properties:

- Session timeout in seconds (data\_analytics\_studio\_session\_timeout) under Ambari Data Analytics Studio CONFIGS Advanced data\_analytics\_studio-webapp-properties
- knoxsso.token.ttl under Ambari Knox CONFIG Advanced knoxsso-topology

The data\_analytics\_studio\_session\_timeout property is a DAS-specific configuration which is used to define the time period for which you want DAS to keep you logged in. This, particularly, comes in handy when you have a long

running query. If a user is running a query which runs for a long time, say more than 24 hours, and if the user gets logged out before the query execution finishes, the user may lose the query results. To prevent this issue, you (the DAS admin) must configure the time period in the Session timeout in seconds field accordingly. By default, DAS does not log you out for 24 hours or 86400 seconds after you log in.

The `knoxsso.token.ttl` property is cluster-wide configuration which applies to all the services in the cluster that use Knox SSO. The logout time is given in milliseconds.

If the user is actively accessing DAS, then he will not be logged out. However, if the user is inactive, then the session timeout is equal to the value specified in either the Session timeout in seconds field in Ambari or in the `knoxsso.token.ttl` property - whichever is higher.

For example, if you have set 5 minutes in the `knoxsso.token.ttl` property and 86400 seconds (24 hours) in the Session timeout in seconds field, then the user will remain logged into DAS for 24 hours, or until the user is actively using DAS.

## Getting a 401 - Unauthorized access error message while accessing DAS

You get an “Unauthorized access” error when you click Data Analytics Studio UI from Ambari Data Analytics Studio Quick Links . This can happen in case your cluster is secured with Knox proxy.

To fix this issue, you must replace the contents of the `quicklink.json` file which is present in the `/var/lib/ambari-server/resources/mpacks/data-analytics-studio-mpack-<VERSION>/hdp-addon-services/DATA_ANALYTICS_STUDIO/<VERSION>/quicklinks/` directory with the template provided in [Setting up quick links for the DAS UI](#).

### Setting up quick links for the DAS UI

To use the DAS UI quick link from Ambari in case of Knox proxy clusters, you (the admin user) need to paste the following `quicklink.json` file in the Ambari server machine.

#### Procedure

1. Create the `quicklink.json` file as per the following template:

```
{
  "name": "default",
  "description": "default quick links configuration",
  "configuration": {
    "protocol": {
      "type": "HTTPS_ONLY"
    },
    "links": [
      {
        "name": "data_analytics_studio_ui",
        "label": "Data Analytics Studio UI",
        "requires_user_name": "false",
        "component_name": "KNOX_GATEWAY",
        "url": "%@://%@:%@/gateway/<DAS_UI_TOPOLOGY_NAME>/das/",
        "port": {
          "https_property": "gateway.port",
          "https_default_port": "8443",
          "regex": "^(\d+)$",
          "site": "gateway-site"
        }
      }
    ]
  }
}
```

Replace the `DAS_UI_TOPOLOGY_NAME` with the actual DAS UI topology.

2. Change to the following directory:

```
/var/lib/ambari-server/resources/mpacks/data-analytics-studio-mpack-<VERSION>/hdp-addon-services/  
DATA_ANALYTICS_STUDIO/<VERSION>/quicklinks/
```

Replace the <VERSION> with the version of the DAS distribution that you have installed.

3. Replace the existing quicklink.json file with the file that you created in step 1.  
Verify that you have correctly replaced the quicklink.json file.
4. Restart Ambari by using the following command:

```
ambari-server restart
```

## Installing DAS on CDP Private Cloud Base using Cloudera Manager

You can install Data Analytics Studio (DAS) using Cloudera Manager if you have transitioned from an Ambari-managed CDP Private Cloud Base cluster to a Cloudera Manager-managed cluster using the AM2CM tool, or if you were using Ambari Views on HDP 2.6.5 or DAS on HDP 3.1.5 and have upgraded to CDP Private Cloud Base 7.x.

### Before you begin



**Attention:** This task is applicable only if you are upgrading to CDP Private Cloud Base 7.1.8 or lower. Starting with the 7.1.9 release, Hue replaces DAS.

Make sure that you have created a PostgreSQL database version 9.6, 10, or 12 to use with DAS. If you have an existing Postgres database that you used with DAS on HDP 3.1.5, then you can point DAS to the same database during the database setup stage, so that you do not lose saved queries and query history.

### Procedure

1. Log in to Cloudera Manager as an Administrator.
2. From the Cloudera Manager homepage, go to the Status tab, click the more options icon on your \$Cluster Name click Add Service.
3. On the **Select Services** page, select Data Analytics Studio and click Continue.
4. On the **Select Dependencies** page, select the dependencies that you want to set up with DAS and click Next.
5. On the **Assign Roles** page, select role assignments for your dependencies and click Continue.
6. On the **Setup Database** page, specify the Postgres database name, database host, username, and password. Click Test Connection to validate the settings and click Continue.

If the connection is successful, a green checkmark and the word Successful appears next to the service name. If there are any problems, an error is reported indicating that the service failed to connect.

7. The **Review Changes** page lists default and suggested settings for several configuration parameters, including data directories.

Select user authentication mode as None.



**Note:** SPNEGO as the default mode of authentication. If you do not want to use SPNEGO, then select None as the user authentication mode.

Review and make any necessary changes, and then click Continue. The **Command Details** page is displayed.

8. The **Command Details** page shows the status of your operation. After completion, you can view the installation logs under the standard output (stdout).

After the First Run command completes, click Continue to go to the **Summary** page.

9. The **Summary** page reports the success or failure of the setup wizard. Click Finish to complete the wizard. The installation is complete.

If the DAS WebApp fails to start with the following error message, then specify the latest activated CDH 7.1.x parcel directory in the DAS WebApp Advanced Configuration Snippet:

```
Error: Could not find or load main class org.apache.ranger.credentialapi
.buildks.
```

For more information, see [Unable to start DAS after installing the GPL Extras parcel](#).

### What to do next

Cloudera recommends that you change the default password as soon as possible by clicking the logged-in username at the top right of the home screen and clicking Change Password.

## Adding Hue service with Cloudera Manager

Hue is a web-based interactive query editor that enables you to interact with databases and data warehouses. Data architects, SQL developers, and data engineers can use Hue to create data models, clean data to prepare it for analysis, and to build and test SQL scripts for applications.

To add the Hue service with Cloudera Manager, complete the following steps:

1. [Install and configure MySQL database](#) on page 317
2. [Add the Hue service using Cloudera Manager](#) on page 317
3. [Enable Kerberos for authentication](#) on page 319
4. [Integrate Hue with Knox](#) on page 320
5. [Grant Ranger permissions to new users or groups](#) on page 322

### Install and configure MySQL database

Hue requires a SQL database such as MySQL, PostgreSQL, or Oracle to store small amounts of data, including user account information, query history, and permissions. The instructions in this topic are for installing the MySQL database.

#### Before you begin

Install MySQL on a host within your cluster. For instructions, refer to the [MySQL Installation Guide](#).

#### About this task

To configure Hue to store data in the MySQL database:

#### Procedure

1. SSH into the host on which you have installed MySQL.
2. Create a new database in MySQL as follows:

```
mysql> create database hue;
```

3. Run the following command to grant privileges to a Hue user to manage the database:

```
mysql> grant all on hue.* to 'hue'@'localhost' identified by 'secretpass
word';
```

### Add the Hue service using Cloudera Manager

After you have installed and configured the MySQL database for Hue, add the Hue service using Cloudera Manager.

## Procedure

1. Sign in to Cloudera Manager as an Administrator.
2. Go to **Clusters Actions** and click **Add Service**.  
The **Add Service to Cluster** wizard is displayed.
3. Select **Hue** as the Service Type and click **Continue**.  
The **Select Dependencies** page is displayed.
4. On the **Select Dependencies** page, select the optional dependencies that are required for your cluster and click **Continue**.  
The required dependency is selected by default.



### Note:

In CDH 6 and earlier, the Hive service included the Hive Metastore and HiveServer2. In Cloudera Runtime 7.0 and later, this service only includes the Hive Metastore. HiveServer2 and other components of the Hive execution engines are part of the HIVE\_ON\_TEZ service.

If you need to run Hive queries, then select an optional dependency containing the HIVE\_ON\_TEZ service.

If you need to run only Impala queries, then the HIVE\_ON\_TEZ service is not required.

5. On the **Assign Roles** page, you can perform one of the following actions:
  - To accept the default role assignments, click **Continue** in the lower right corner of the page. To view the default role assignments by host, click **View By Host**.
  - To customize role assignments, click the field below Hue Server and Load Balancer roles. This launches a dialog box where you can select hosts where you want to add the role. Click **OK** after making your selection.

The wizard evaluates host hardware configurations to determine the best hosts for each role. All worker roles are automatically assigned to the same set of hosts as the HDFS DataNode. You can reassign if necessary. Specify hostnames by IP address, rack name, or by range:

| Range Definition        | Matching Hosts   |
|-------------------------|--|
| 10.1.1.[1-4]            | 10.1.1.1, 10.1.1.2, 10.1.1.3, 10.1.1.4   |
| host[1-3].company.com   | host1.company.com, host2.company.com, host3.company.com                        |
| host[07-10].company.com | host07.company.com, host08.company.com, host09.company.com, host10.company.com |

6. On the **Setup Database** page:
  - a) Select a database vendor from the **Type** field and specify the Database Hostname, Database Name, Username, and Password.
  - b) Click **Test Connection**, and when the success message appears, click **Continue** and Cloudera Manager starts the Hue service.
7. Review your configurations on the **Review Changes** page and click **Continue**.
8. On the **Command Details** page, expand the running commands to view the details of any step, including log files and command output. You can filter the view by selecting **Show All Steps**, **Show Only Failed Steps**, or **Show Only Running Steps**.  
After the First Run command completes, click **Continue** to go to the **Summary** page.
9. After the service is started, click **Continue** and then click **Finish**.  
If your cluster uses Kerberos, Cloudera Manager automatically adds a Hue Kerberos Ticket Renewer role to each host where you assigned the Hue Server role instance. See "Enabling Kerberos Authentication" for more information.

## Enable Kerberos for authentication

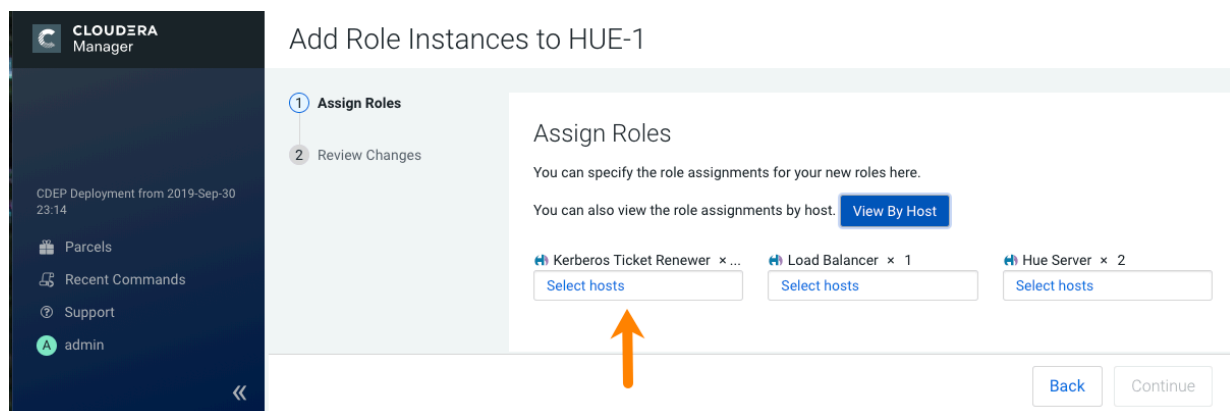
For Hue to work properly with a CDP cluster that uses Kerberos for authentication, the Kerberos Ticket Renewer role must be added to the Hue service. After you have added the Hue service to your cluster, renew the Kerberos ticket for Hue and enable Kerberos for authentication.

### About this task

Use the Cloudera Manager Admin Console to add the Kerberos Ticket Renewer role to each host with a Hue Server role instance. The Hue Kerberos Ticket Renewer renews only those tickets created for the Hue service principal: `hue/HOSTNAME@REALM-NAME`. The Hue principal impersonates other users for applications within Hue such as the Job Browser, File Browser, and so on. Other services, such as HDFS and MapReduce, do not use the Hue Kerberos Ticket Renewer. Instead these other services handle ticket renewal as needed by using their own mechanisms.

### Procedure

1. On the Cloudera Manager home page, select the Hue service.
2. On the Hue service page, click the Instances tab.
3. On the Instances page, click Add Role Instances on the right side of the page. This launches the Add Role Instances wizard.
4. To add a Kerberos Ticket Renewer role instance to the same host that has the Hue server on your CDP cluster, click Select hosts under Kerberos Ticket Renewer:



To check which host has the Hue Server role instance, click View By Host, which launches a table that lists all the hosts in your CDP cluster and shows all the roles each host already has.

5. In the host selection dialog box, after selecting the host where you want to add the Kerberos Ticket Renewer role instance, click OK, and Cloudera Manager adds the role instance.
6. After processing the request to add the role instance, Cloudera Manager returns you to the Instances page and prompts you to restart the service. Click the Restart the service (or the instance)... link so the configuration change can take effect.
7. After the services have restarted, click Finish to return to the Instances page.

Repeat these steps for each Hue Server role on your cluster.

### What to do next

Troubleshooting the Kerberos Ticket Renewer:

If the Hue Kerberos Ticket Renewer does not start, check the configuration of your Kerberos Key Distribution Center (KDC). Look at the ticket renewal property, `maxrenewlife`, to ensure that the principals, `hue/<HOST_NAME>` and `krbtgt`, are renewable. If these principals are not renewable, run the following commands on the KDC to enable them:

```
kadmin.local: modprinc -maxrenewlife 90day krbtgt/<YOUR_REALM.COM>
```

```
kadmin.local: modprinc -maxrenewlife 90day +allow_renewable hue/  
<HOST_NAME>@<YOUR_REALM>
```

## Integrate Hue with Knox

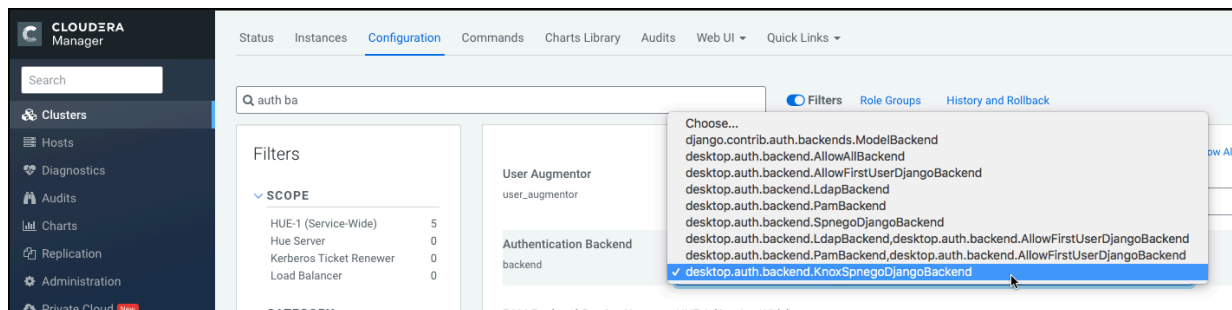
You can use the Apache Knox Gateway to interact with Hue REST APIs and the Hue user interface, along with other CDP components and services. To set up Knox Single Sign-on (SSO) to authenticate users, you must configure the KnoxSpnegoDjangoBackend property using Cloudera Manager.

### Before you begin

To authenticate users using Knox SSO, you must have Knox installed on your CDP cluster, also known as a secure cluster.

### Procedure

1. Sign in to Cloudera Manager as an Administrator.
2. Go to **Clusters**  $\$$ Hue service **Configurations** and search for the **Authentication Backend** field.
3. Select `desktop.auth.backend.KnoxSpnegoDjangoBackend` from the dropdown.




4. Click **Save Changes**.
5. Go to **Clusters**  $\$$ Knox service **Instances** and note down the hostnames of the Knox Gateways.

You must provide these details in the next step.

If you have set up Knox in High-Availability (HA) mode, then you can see more than one Knox Gateways listed on the **Instances** tab.

6. Go back to **Clusters**  $\$$ Hue service **Configurations** and search for the **Knox Proxy Hosts** field.
7. Enter the hostname of the Knox Gateway that you noted earlier.

If you have set up Knox HA, then click **+** to add another hostname.



8. If you have deployed a Hue Load Balancer, then you must specify the Load Balancer hostname in the **Knox Proxy Hosts** field by clicking **+**.
9. Click **Save Changes**.

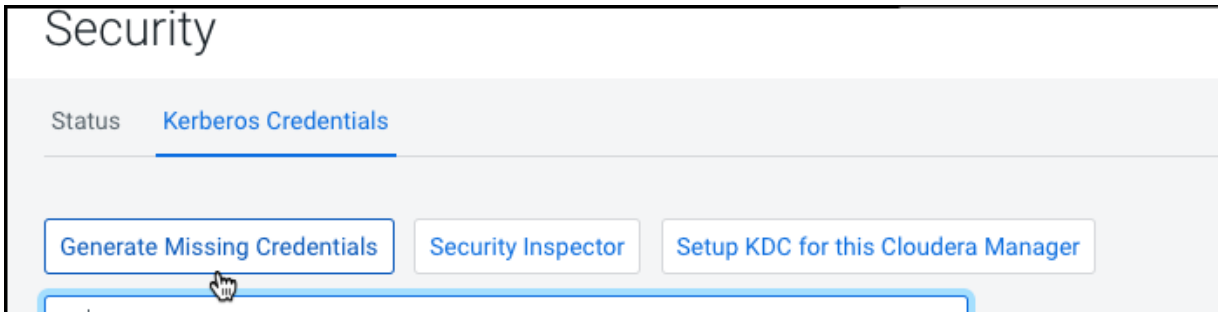
You would see the following warning:

Role is missing Kerberos keytab. Go to the **Kerberos Credentials** page and click the **Generate Missing Credentials** button.

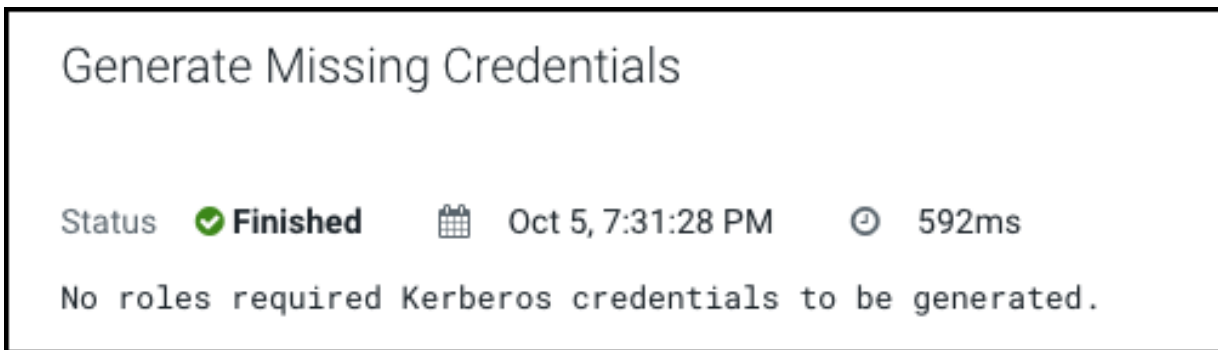
10. Click **Administration** on the Cloudera Manager left navigation panel and select **Security**.




11. Go to the Kerberos Credentials tab and click Generate Missing Credentials.



A pop-up showing the status is displayed.



12. Go to Clusters \$Hue service and click Restart next to Actions.



13. On the Stale Configurations page, click Restart Stale Services.


The Restart Stale Services wizard is displayed.

14. On the Review Changes page, select Redeploy client configuration, and click Restart Now.

The Command Details page shows the live status as the service restarts.

When all the steps are complete, click Finish.

15. From the Hue service page, click Web UI Knox Gateway UI .

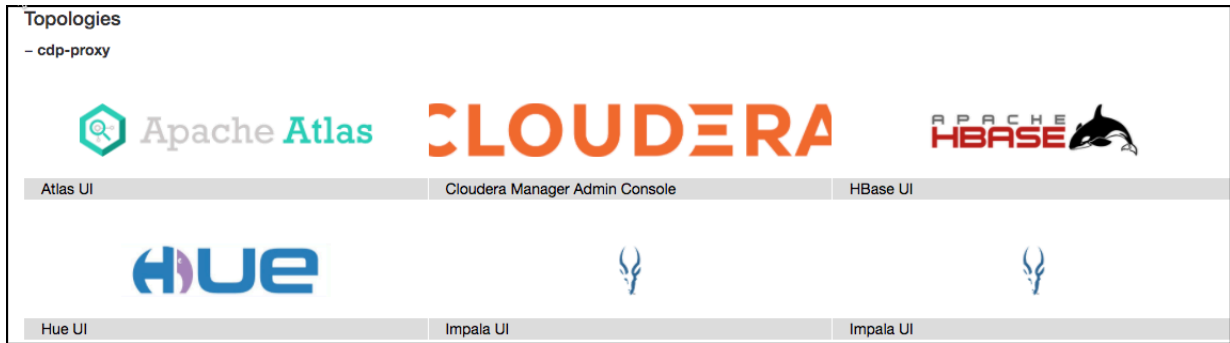


The Knox Gateway UI is displayed.

16. On the General Proxy Information page, expand the CDP Proxy topology by clicking + cdp-proxy under Topologies.

The list of services that are configured with the cdp-proxy topology is displayed.

17. Click on the Hue logo.



You should be able to log in to the Hue web UI.

You can also log into Hue using the following URL:

```
https://[***HOSTNAME***]:[***PORT***]/gateway/cdp-proxy/hue/
```

## Grant Ranger permissions to new users or groups

If you have secured your cluster using Ranger, then you must grant proper permissions to your users and groups from the Ranger web interface. If your users do not have proper permissions, then they may not be able to access certain databases or tables from the Hue editor.

### About this task

To grant permissions to a new user or group:

### Procedure

1. Log into Cloudera Manager as an Administrator.
2. Go to Clusters \$Ranger service Instances tab and note down the hostname corresponding to the “Ranger Usersync” role type.
3. Open the Ranger UI by clicking Ranger Admin Web UI.
4. SSH into the Ranger Usersync host that you noted in step 2 and add the user or the group as follows:

```
ssh root@example.domain.site useradd [***USERNAME/GROUP-NAME***] passwd [***PASSWORD***]
```

5. On the Ranger web UI, click Hadoop SQL listed under the HADOOP SQL service. The **Hadoop SQL Policies** page is displayed.

6. On the **Hadoop SQL Policies** page, you can grant the new user access to all the databases or to specific databases by adding a new policy.

To grant the permission on all databases:

- a. Click the policy ID corresponding to "all - database, table, column".

| Policy ID | Policy Name                   | Policy Labels | Status  | Audit Logging | Roles | Groups | Users                               | Action                |
|-----------|-------------------------------|---------------|---------|---------------|-------|--------|-------------------------------------|-----------------------|
| 7         | all - global                  | --            | Enabled | Enabled       | --    | --     | hive beacon dpprofiler hue + More.. | [Eye] [Edit] [Delete] |
| 8         | all - database, table, column | --            | Enabled | Enabled       | --    | --     | hive beacon dpprofiler hue + More.. | [Eye] [Edit] [Delete] |
| 9         | all - database, table         | --            | Enabled | Enabled       | --    | --     | hive beacon dpprofiler hue + More.. | [Eye] [Edit] [Delete] |
| 10        | all - database                | --            | Enabled | Enabled       | --    | public | hive beacon dpprofiler hue + More.. | [Eye] [Edit] [Delete] |
| 11        | all - hiveservice             | --            | Enabled | Enabled       | --    | --     | hive beacon dpprofiler hue + More.. | [Eye] [Edit] [Delete] |

- b. On the **Edit Policy** page, add the user whom you want to grant the permission in the Select User field under the Allow Conditions section as shown in the following image:

| Select Role  | Select Group  | Select User   | Permissions   | Delegate Admin                                     |
|--------------|---------------|---|---|--|
| Select Roles | Select Groups | <input type="checkbox"/> hive <input type="checkbox"/> beacon <input checked="" type="checkbox"/> dpprofiler <input type="checkbox"/> hue <input type="checkbox"/> admin <input checked="" type="checkbox"/> impala | <input checked="" type="checkbox"/> select <input checked="" type="checkbox"/> update <input type="checkbox"/> Create <input type="checkbox"/> Drop <input type="checkbox"/> Alter <input type="checkbox"/> Index <input type="checkbox"/> Lock <input type="checkbox"/> All <input type="checkbox"/> Read <input type="checkbox"/> Write <input type="checkbox"/> ReplAdmin <input type="checkbox"/> Service Admin <input type="checkbox"/> Temporary UDF Admin <input type="checkbox"/> Refresh | <input checked="" type="checkbox"/> Delegate Admin |
| Select Roles | Select Groups | <input type="checkbox"/> rangerlookup   | <input type="checkbox"/> Read   | <input type="checkbox"/> Delegate Admin            |
| Select Roles | Select Groups | <input type="checkbox"/> (OWNER)  | <input type="checkbox"/> All  | <input checked="" type="checkbox"/> Delegate Admin |

To grant permissions to a group, enter the group name in the Select Group field.

- c. Click Save.

To grant permission on specific database:

- a. Click Add New Policy.

The **Create Policy** page is displayed.

- b. Under the Policy Details section, specify the policy name and select the database, table, and column that you want your user to access as shown in the following image:

**Policy Details :**

Policy Type: Access Add Validity Period

Policy Name: default database, tables, column enabled no

Policy Label: Policy Label

database: default include

table: \* include

column: \* include

Description: This policy grants users permissions to only the "default" database and its tables and columns.

Audit Logging: YES

**Allow Conditions :**

| Select Role  | Select Group  | Select User  | Permissions   | Delegate Admin                          |
|--------------|---------------|--------------|---|---|
| Select Roles | Select Groups | Select Users | <input checked="" type="checkbox"/> select <input checked="" type="checkbox"/> update <input checked="" type="checkbox"/> Create <input checked="" type="checkbox"/> Drop <input type="checkbox"/> Alter <input type="checkbox"/> Index <input type="checkbox"/> Lock <input type="checkbox"/> All <input checked="" type="checkbox"/> Read <input checked="" type="checkbox"/> Write <input type="checkbox"/> ReplAdmin <input type="checkbox"/> Service Admin <input type="checkbox"/> Temporary UDF Admin <input checked="" type="checkbox"/> Refresh <input type="checkbox"/> Select/Deselect All | <input type="checkbox"/> Delegate Admin |

- c. Under the Allow Conditions section, enter the username in the Select User field and click Add Permissions and select the permissions that your user must have.

To grant permissions to a group, enter the group name in the Select Group field.

- d. Click Add.

7. Start the Hue service from Cloudera Manager.

### Results

The user or the group should be able to run any query on any entities as defined in the policy.

## Adding Query Processor service to a cluster

The Query Processor service indexes Hive and Tez events and provides APIs to access them. It is required if you want to view the Queries tab (query history and query details) on the Hue Job Browser. You must install the Query Processor service on your clusters manually.

### About this task



**Attention:** This task is applicable only if you are upgrading to 7.1.8 and higher. If you are upgrading to 7.1.7 SP2, 7.1.7 SP1, 7.1.7, or lower, then you can skip this task.



**Note:** The Query Process service supports only MySQL.

You can either install the Query Processor service as a dependency while adding the Hue service or after adding the Hue service.

### Before you begin

This task assumes that you already have a database installed on a host in your cluster. On , Hue Query Processor supports PostgreSQL, MySQL, and Oracle databases. The supported PostgreSQL database version for Hue Query Processor is 9.6 and higher.

Next, you need to create a database for the Query Processor service with the required roles. To create the Query Processor database:


1. SSH into your database host as a root user.
2. Start the psql terminal by entering the following command:

```
sudo -u postgres psql
```

3. Run the following statement while specifying the username, password, and a database name for the Query Processor:

```
CREATE ROLE [***QP-USER***] WITH LOGIN PASSWORD '[***QP-PASSWORD***]';
ALTER ROLE [***QP-USER***] CREATEDB;
CREATE DATABASE [***QP-DATABASE***];
GRANT ALL PRIVILEGES ON DATABASE [***QP-DATABASE***] TO [***QP-USER***];
```

### Procedure

1. Log in to as an Administrator.
2. Go to Clusters  Add Service .
3. Select Query Processor on the **Add Service to Cluster** page and click Continue.  
The required dependencies are automatically selected.

4. Select the host on which you want to install the Query Processor by clicking View By Host. Then click Continue.
5. Select the database type, and specify the database hostname, database name, and username and password to access the database on the **Setup Database** page and click Test Connection.  
After the connection is verified, click Continue.
6. On the **Review Changes** page, accept the default settings and click Continue.  
The **Command Details** page is displayed, which shows the status of the installation and configuration.
7. Click Continue if the operation is successful.
8. Click Finish on the **Summary** page.

### Results

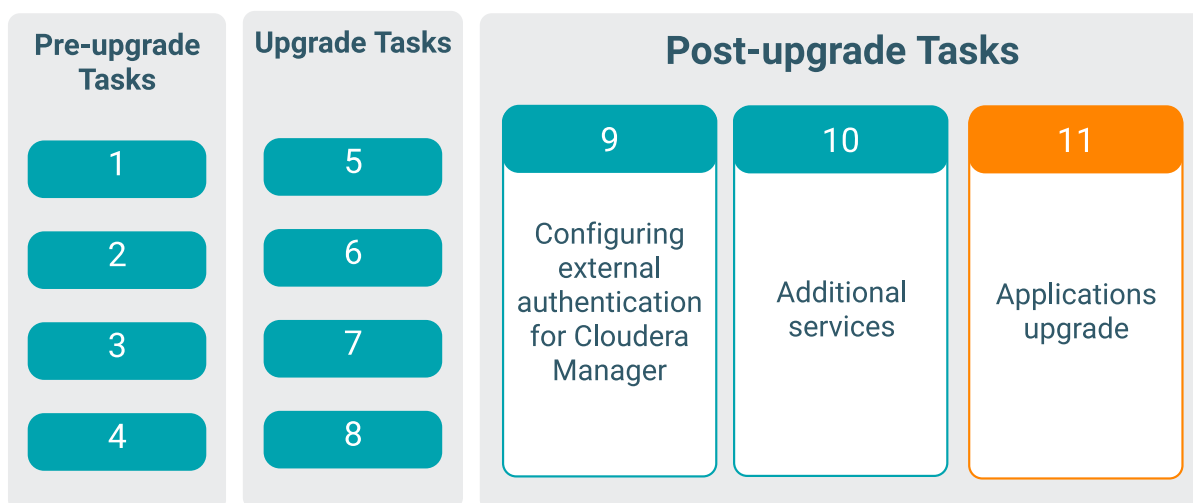
You are redirected to the home page. You can now see the Query Processor service listed within your cluster.

### What to do next

(Optional) To view the **Queries** tab on the **Job Browser** page in Hue, you must enable the Query Processor Service option in Hue configurations using .

## Applications Upgrade

After you upgrade, you must test all the services that run on your platform.



Ideally, you should have an indicative subset of jobs from your workloads. These are the tasks that you should have identified and run before the upgrade allowing you to compare pre-upgrade versus post-upgrade test results. These tests should also include any parts of the application that required code changes due to the changes in the platform. For example, to cater for changes in Hive managed versus external tables. The tests should also include a performance test. This can help to highlight missed or wrong configuration settings or point to other issues with the upgrade. Depending on your environment, perform these steps.

1. Update application code with changes required by the upgraded platform
2. Update the dependencies in the pom.xml file for custom jar files used in your applications to use the new dependencies for CDP Private Cloud Base.
3. Restart applications
4. Test the applications and verify they are functioning and performing as they were prior to upgrade



**Note:** After successfully upgrading from HDP to CDP Private Cloud Base, you can remove the HDP bits from the CDP Private Cloud Base cluster using the yum commands. Otherwise, when you run any security tool or security scanner for CVE detection, the HDP bits on the CDP Private Cloud Base cluster are detected as CVEs.

## Procedure to Rollback from CDP Private Cloud Base 7.1.7 SP1 to CDP Private Cloud Base 7.1.7

You can rollback from the CDP Private Cloud Base 7.1.7 SP1 to CDP Private Cloud Base 7.1.7. To rollback to CDP Private Cloud Base 7.1.7:

### Procedure

1. Log in to the Admin Console.
2. Click Parcels from the left menu.
3. Click Parcel Repositories & Network Settings.



**Note:** If the 7.1.7 URL is available in the parcel repository, then you can skip to step 8. If not, then proceed with step 4.

4. In the Remote Parcel Repository URLs section, click the "+" icon and add the 7.1.7 URL for your Parcel repository.
5. Click Save & Verify Configuration. A message with the status of the verification appears above the Remote Parcel Repository URLs section. If the URL is not valid, check the URL and enter the correct URL.
6. After the URL is verified, click Close.
7. Locate the row in the table that contains the new parcel and click the Download button.
8. After the download of the new parcel is complete, click the Distribute button.  
Wait for the parcel to be distributed and unpacked before continuing. displays the status of the parcel distribution. Click on the status display to view detailed status for each host.
9. Click Activate. parcels are activated on the cluster.
10. When the parcel is activated, click the logo to return to the home page.  
The cluster is now rolled back to the CDP Private Cloud Base 7.1.7 version.
11. When the parcel is activated, click the Actions menu next to the cluster name and select Post Cloudera Runtime Upgrade.
12. Restart the cluster: Click the Actions menu and select Restart.