Cloudera Manager 7.11.3

Release Notes

Date published: 2020-11-30 Date modified: 2024-02-23



https://docs.cloudera.com/

Legal Notice

© Cloudera Inc. 2024. All rights reserved.

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

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

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

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

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

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

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

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

Contents

What's New in Cloudera Manager 7.11.3	
What's new in Platform Support	
Known Issues in Cloudera Manager 7.11.3	
Fixed Issues in Cloudera Manager 7.11.3	
Fixed Common Vulnerabilities and Exposures	
Deprecation notices in Cloudera Manager 7.11.3	
Deprecation Notices for Cloudera Manager	
Platform and OS	
Cumulative hotfixes	
Cloudera Manager 7.11.3 Cumulative hotfix 5	
Cloudera Manager 7.11.3 Cumulative hotfix 4	
Cloudera Manager 7.11.3 Cumulative hotfix 3	
Cloudera Manager 7.11.3 Cumulative hotfix 2	
Cloudera Manager 7.11.3 Cumulative hotfix 1	

What's New in Cloudera Manager 7.11.3

New features and changed behavior for Cloudera Manager 7.11.3.

New features

Zero Downtime Upgrades (ZDU)

Cloudera is reintroducing the concept of rolling upgrades in CDP 7.1.9 in an easier to use format called Zero Downtime Upgrades (ZDU). Zero Downtime Upgrades automates the process of performing rolling upgrades in an optimized format to allow for minimal to zero downtime depending on the services installed on a cluster. All future service packs and runtime upgrades will support ZDU. However, the enhancements brought by ZDU will be available on upgrades from CDP 7.1.7 and CDP 7.1.8. Before using this feature read the upgrade instructions. For more information, see the Zero Downtime upgrade documentation.



Caution: Clouder recommends all upgrades to happen in a maintenance window by throttling and scaling down workloads during that time as a best practice.

TLS 1.2 encryption support for secured database connections

Cloudera Manager supports TLS (Transport Layer Security) 1.2 encryption between the Cloudera Manager Server and the backend databases such as MySQL, PostgreSQL, and MariaDB.

Now you can enable TLS 1.2 on Database Server and Cloudera Manager Server in the database environment. See Configuring TLS 1.2 for Cloudera Manager.

Also, now you can enable TLS 1.2 on Reports Manager in the database environment. See Configuring TLS 1.2 for Reports Manager.

Cloudera recommends that you secure the network connection between the Cloudera Manager Server and the backend database using TLS 1.2 encryption.

The scm_prepare_database.sh script in Cloudera Manager now accepts the following two new optional parameters:

- -s|--ssl
- --jdbc-url

For more information on optional parameters, see Syntax for scm_prepare_database.sh.

TCPS support for connections to Oracle database

Cloudera Manager supports connections to backend Oracle database that are secured with Transmission Control Protocol with SSL (TCPS). This provides greater security for connections between Cloudera Manager Server and the backend Oracle database. For more information, see Enabling TCPS for Oracle Database Server.

Now you can enable TCPS on Reports Manager in the Oracle database environment. For more information, see Configuring TCPS for Reports Manager.

Python 3.8 (or 3.9 for RHEL 9.1) support for Cloudera Manager 7.11.3

Cloudera Manager 7.11.3 requires Python 3.8 on most of the supported operating systems. The exception is that on the RHEL 9.1 operating system, it supports Python 3.9 version only.

Cloudera Manager 7.11.3 does not work with Python 2.7. While using Cloudera Manager 7.11.3 with Cloudera Runtime 7.1.8 or 7.1.9 version, you may remove all Python 2 versions from the operating system, only when the operating system allows you to remove the Python 2 version.

If you are running Cloudera Runtime 7.1.7 SP2 or below versions with Cloudera Manager 7.11.3, then Python 2.7 is still required for the Cloudera Runtime components. In this scenario, you must

install both Python 2.7 (for Cloudera Runtime components) and Python 3 (for Cloudera Manager 7.11.3).

You must install Python 3.8 (or 3.9 for RHEL 9.1) on all hosts before upgrading to Cloudera Manager 7.11.3. See Installing Python 3.

For more information about the operating systems that are supported when using Python 3.x with the Cloudera Manager Agents, see Platform support for Cloudera Manager 7.11.3.

Support for noexec option on the /tmp directory

Cloudera Manager functions normally when you enable the noexec option for the /tmp directory on cluster hosts.

The /tmp directory on Linux hosts is used by many applications to store non-persistent data and to execute transient scripts.

Users require this noexec option on /tmp directory to eliminate possible security risks by preventing the execution of binaries from the /tmp filesystem.

The noexec option prevents unintentional system modifications or corruption that may potentially lead to system instability or data theft.

Ability to modify existing Data Context and allow Ozone to be an option

Data Contexts in Cloudera Manager are used to access data in Cloudera Private Cloud Base environment. You can add or remove certain services to the Data Context. See About Data Context and Creating a Compute Cluster and Data Context.

Certify CM with HA Postgres databases with SSL enabled

Postgres HA support involves enabling Postgres HA and configuring Postgres HA behind a load balancer. See PostgresSQL High Availability.

Iceberg replication policies

You can create Iceberg replication policies in CDP Private Cloud Base Replication Manager to replicate Iceberg tables between CDP Private Cloud Base 7.1.9 or higher clusters using Cloudera Manager 7.11.3 or higher versions.

For more information, see Iceberg replication policies

Ranger replication policies

You can create Ranger replication policies in CDP Private Cloud Base Replication Manager. The Ranger replication policies migrate Ranger policies for HDFS, Hive, and HBase services between Kerberos-enabled CDP Private Cloud Base 7.1.9 or higher clusters using Cloudera Manager 7.11.3 or higher versions.

For more information, see Ranger replication policies.

Incremental replication of Ozone data using Ozone replication policies

You can choose the "Full file listing", "Incremental only", or "Incremental with fallback to full file listing" option as a Listing method during the Ozone replication policy creation process. The listing method determines the replication method that Ozone replication policy can use to replicate Ozone data.

For more information, see Ozone replication policies.

Ozone snapshot policies

You can create Ozone snapshot policies in CDP Private Cloud Base Replication Manager to take snapshots of Ozone buckets and volumes at regular intervals. Ozone replication policies leverage the snapshots to perform incremental replication. You can also restore an Ozone bucket to an earlier version using snapshots or restore the Ozone bucket to another bucket in Cloudera Manager.

For more information, see Ozone snapshot policies.

Collecting Heartbeat data from Cloudera Manager

Beginning with Cloudera Manager 7.11.3, a report containing basic cluster information will securely transmit to Cloudera periodically. This report contains cluster-related metadata to determine the version and size of each cluster. This information will assist Cloudera in gaining a clearer understanding of our customers' deployments so we can deliver more robust support and an improved customer experience.

Reports will be saved locally for Customers with infrastructure isolated from the public internet. For assistance, please open a General Administrative Assistance case on MyCloudera.

The generated report is human-readable for users and can be found under /var/lib/cloudera-scm-server/reports (configured as default).

Replicate Hive external tables in Dell EMC Isilon storage clusters using Hive external table replication policies

You can use Hive external table replication policies in CDP Private Cloud Base Replication Manager to replicate Hive external tables between Dell EMC Isilon storage clusters where the 7.1.9 clusters use Cloudera Manager 7.11.3 CHF1 or higher versions.

Changed or updated features

An UI for Credential Storage Provider (CSP) is introduced on Cloudera Manager interface

On Cloudera Manager UI, now you can enable and manage CSP. To find CSP, go to Administration Security Status tab.

From this release onwards, CSP is generally available (GA). CSP is used to encrypt the sensitive values by configuring a Secure Credential Store that stores an encryption key to encrypt and decrypt sensitive information. Later this sensitive information is stored in encrypted form only in the Cloudera Manager database. For more information about CSP, see Securing sensitive information using a Secure Credential Storage Provider.

Remove the SHA-1 hashing algorithm based GPG signing key and update them with the SHA-256 based GPG key

Cloudera Manager install packages (RPM and Deb) are now signed with the SHA-256 hashing algorithm. You must remove the SHA-1 hashing algorithm based GPG signing key and update them with the SHA-256 based GPG key.

From this release onwards, you must import a new GPG public key into the OS key ring when installing the Cloudera Manager Agent, Cloudera Manager Server, and Cloudera Manager Daemon packages. The SHA-256 based signing key is applicable to both the fresh installation and upgrade of Cloudera Manager (7.11.3 version). The new GPG keys are now signed with a more secure SHA-256 hashing algorithm.



Important: You must incorporate the new RPM-GPG-KEY-cloudera into your installation or upgrade scripts if you use any automation scripts to install or upgrade Cloudera Manager.

Note:

The new GPG key for operating systems such as RHEL 9, RHEL 8, RHEL 7, and SLES 15 is located here.

For Debian based operating systems such as Ubuntu18 and Ubuntu 20, the new GPG key is located here.

Platform support for Cloudera Manager 7.11.3

The following table provides the details about the operating systems that are supported when using Python 3.x with the Cloudera Manager Agents:

Python 3.8	Python 3.9
• RHEL 7	• RHEL 9
• RHEL 8	
Oracle 8.8 UEK	
• SLES 12	
• SLES 15	
• Ubuntu 20	

For more information about the minor version operating system support, see Cloudera Support Matrix.

What's new in Platform Support

You must be aware of the platform support for the Cloudera Manager 7.11.3 release.

Platform Support Enhancements

- **New OS support**: Cloudera Manager 7.11.3 now supports the following operating systems:
 - RHEL 9
 - RHEL 8
 - RHEL 8.8 FIPS (added RHEL 8.8 support for FIPS customers with JDK8)
 - Oracle 8.8 UEK
 - SLES 15 SP4

For more information about the minor version operating system support, see Cloudera Support Matrix.

Known Issues in Cloudera Manager 7.11.3

Known issues in Cloudera Manager 7.11.3.

OPSAPS-68689: Unable to emit the LDAP Bind password in core-site.xml for client configurations

If the CDP cluster has LDAP group to OS group mapping enabled, then applications running in Spark or Yarn would fail to authenticate to the LDAP server when trying to use the LDAP bind account during the LDAP group search.

This is because the LDAP bind password was not passed to the /etc/hadoop/conf/core-site.xml file. This was intended behavior to prevent leaking the LDAP bind password in a clear text field.

Set the LDAP Bind password through the HDFS client configuration safety valve.

- **1.** On the Cloudera Manager UI, navigate to the HDFS service, by clicking on the HDFS service under the Cluster.
- **2.** Click the Configuration tab. Search for the HDFS Client Advanced Configuration Snippet (Safety Valve) for hdfs-site.xml configuration parameter.
- **3.** Add an entry with the following values:
 - Name = hadoop.security.group.mapping.ldap.bind.password
 - Value = (Enter the LDAP bind password here)
 - Description = Password for LDAP bind account
- 4. Then click the Save Changes button to save the safety valve entry.
- **5.** Perform the instructions from the Manually Redeploying Client Configuration Files to manually deploy client configuration files to the cluster.

OPSAPS-69342: Access issues identified in MariaDB 10.6 were causing discrepancies in High Availability (HA) mode

MariaDB 10.6, by default, includes the property require_secure_transport=ON in the configuration file (/etc/my.cnf), which is absent in MariaDB 10.4. This setting prohibits non-TLS connections, leading to access issues. This problem is observed in High Availability (HA) mode, where certain operations may not be using the same connection.

To resolve the issue temporarily, you can either comment out or disable the line require_secure_t ransport in the configuration file located at /etc/my.cnf.

OPSAPS-68577: Invalid Iceberg license validator message

During the Cloudera Manager 7.11.3 (Cloudera Runtime 7.1.9) upgrade, you might see the following warning message:

"details on this warning: Validation Suppress Configuration Validator: Iceberg License Validator Current Message Failed parameter validation. Suppress For CORE_SETTINGS-1"

You can safely suppress the Configuration validator message.

OPSAPS-69357: Python incompatibility issues when Cloudera Manager (Python 3.x compatible) manages a cluster with Cloudera Runtime 7.1.7 (Python 2 compatible)

If Cloudera Manager is compatible with Python 3, then scripts that are packaged with this Cloudera Manager are also ported to Python 3 syntax.

So, using Cloudera Manager (7.11.3 or any other Cloudera Manager version ported to Python 3.x version) to manage a cluster with Cloudera Runtime 7.1.7 (Python 2 compatible) would cause Python incompatibility issues because the process assumes Python 2 environment but the scripts that are packaged with this Cloudera Manager are ported to Python 3 syntax.

None

OPSAPS-68452: Azul Open JDK 8 and 11 are not supported with Cloudera Manager

Azul Open JDK 8 and 11 are not supported with Cloudera Manager. To use Azul Open JDK 8 or 11 for Cloudera Manager RPM/DEBs, you must manually create a symlink between the Zulu JDK installation path and the default JDK path.

After installing Azul Open JDK8 or 11, you must run the following commands on all the hosts in the cluster:

Azul Open JDK 8

RHEL or SLES

```
# sudo ln -s /usr/lib/jvm/java-8-zulu-openjdk-jdk /usr/lib/jvm/j
ava-8-openjdk
```

Ubuntu or Debian

```
# sudo ln -s /usr/lib/jvm/zulu-8-amd64 /usr/lib/jvm/java-8-openj
dk
```

Azul Open JDK 11

For DEBs only

sudo ln -s /usr/lib/jvm/zulu-11-amd64 /usr/lib/jvm/java-11

OPSAPS-69255: Using auth-to-local rules to isolate cluster users is not working consistently

When your cluster is defining auth_to_local rules as mentioned in Using auth-to-local rules to isolate cluster users, after upgrading Cloudera Manager you might experience undesired configuration changes. Many marked Proxyuser settings are removed from core-site.xml and the user 'nobody' is added.

Set Cloudera Manager to the previous behavior for the Proxyuser configurations, this requires editing /etc/default/cloudera-scm-server to add the following JVM argument to the CMF_JAVA_OPTS:

-Dcom.cloudera.cmf.service.config.HadoopUserStrategy=LEGACY

OPSAPS-59723: Extra step required when using Cloudera Manager Trial installer on SLES 15 SP4

When using cloudera-manager-installer.bin to install a trial version of Cloudera Manager, the installation will fail.

Before running cloudera-manager-installer.bin, run the following command:

```
SUSEConnect --list-extensions
SUSEConnect -p sle-module-legacy/15.4/x86_64
zypper install libncurses5
```

OPSAPS-66579: The GUI version of the Cloudera Manager self-installer is not available on the RHEL 9 operating system

While installing the Cloudera Manager (Cloudera Manager Server, Cloudera Manager Agent, and the database) the GUI version of the Cloudera Manager self-installer is not available on the RHEL 9 operating system.

This issue is due to the non-availability of the libncurses5 library on the RHEL 9 operating system. Users can provide input using the CLI prompts instead of the GUI prompts during the installation process.

Use CLI prompts instead of GUI prompts.

OPSAPS-68395: Cloudera Management Service roles might fail to start

While starting Cloudera Manager Server (during a fresh install, an upgrade, or when rolling back an upgrade) the status of one or more roles of the Cloudera Management Service are in the Stopped state, and later these roles might fail to start.

This failure might happen if you attempt to start the affected role(s) within first few minutes after starting the Cloudera Manager Server or a cluster, then the status of the affected roles shows the Down state, and the corresponding functionality is lost. Accordingly, Cloudera Manager might display the errors. This failure is caused by a temporary resource contention, and subsequent timeout.

After fifteen minutes, restart the affected roles, or the Cloudera Management Service as a whole. Alternatively, go to Clusters Cloudera Management Service Configuration and change / increase the value of Descriptor Fetch Max Attempts and Starting Interval for Descriptor Fetch Attempts. Cloudera recommends to set the value of Descriptor Fetch Max Attempts to "30" and Starting Interval for Descriptor Fetch Attempts to "30" seconds.

OPSAPS-60726: Newly saved parcel URL is not showing up on the parcels page in Cloudera Manager High Availability (HA) cluster

Newly saved parcels might not show up on the parcels page in Cloudera Manager HA mode.

You must restart the active and passive Cloudera Manager nodes.

OPSAPS-68178: Inconsistent Java Keystore Type while performing upgrade from CDH 6 to CDP Private Cloud Base 7.1.9

While performing upgrade from CDH 6 to CDP Private Cloud Base 7.1.9, the configured Java Keystore Type is jks on Cloudera Manager UI. However, the physical Truststore files on the upgraded cluster are available in pkcs12 format.

If the value of Java Keystore Type on Cloudera Manager UI is different from the actual Java Keystore Type in the physical Truststore files on the upgraded cluster, then perform the following steps:

1. Stop the Cloudera Manager Server.

sudo systemctl stop cloudera-scm-server

- **2.** Connect to the database.
- 3. Verify the Java Keystore type which is set in database by running the following command:

select * from CONFIGS WHERE ATTR='keystore_type';

- 4. Verify the value of the CONFIG_ID in the result of the previous select command.
- 5. Update the row (previously selected CONFIG_ID) on the database with the correct CONFIG_ID from your cluster by running the following command:

UPDATE CONFIGS SET VALUE ='jks' WHERE CONFIG_ID=config_id;

6. Start the Cloudera Manager Server.

sudo systemctl start cloudera-scm-server

OPSAPS-67929: While upgrading from CDP 7.1.7 SP2 to CDP 7.1.9 version and if there is an upgrade failure in the middle of the process, the Resume option is not available.

You must reach out to Cloudera Support.

OPSAPS-68325: Cloudera Manager fails to install with MariaDB 10.6.15, 10.5.22, and 10.4.31

Cloudera Manager Server fails to execute the DDL commands that involve disabling the FORE IGN_KEY_CHECKS when you use the following databases:

- MariaDB 10.6.15
- MariaDB 10.5.22
- MariaDB 10.4.31

None

OPSAPS-68240: After restarting Cloudera Manager Server and MySQL, Cloudera Manager server fails to start

When using MySQL 8 version, Cloudera Manager fails to start and displays an error message on the logs - java.sql.SQLNonTransientConnectionException: Public Key Retrieval is not allowed



Important: This issue occurs when SSL is enabled on the MySQL side.

To fix this issue, perform the workaround steps as mentioned in the KB article.

If you need any guidance during this process, contact Cloudera support.

DMX-3167

When multiple Iceberg replication policies replicate the same database simultaneously, one of the replication policies might show "Database already exists" error.

Run the replication policy again, the next run for the replication policy succeeds.

DMX-3193

If the source and target clusters have the same nameservice environment and a table is dropped on the source cluster during the incremental replication run of an Iceberg replication policy, the replication policy fails with the "Metadata file not found for table" error. Copy the metadata file from the target cluster to the source cluster and run the incremental replication again.

OPSAPS-68143

When you replicate *empty* OBS buckets using an Ozone replication policy, the policy fails and a FileNotFoundException appears during the "Run File Listing on Peer cluster" step.

DMX-3169

The YARN jobs (DistCp) for Iceberg replication policies cannot use the *hdfs* username if the replication policies use secure source and target clusters.

Provide a proxy user to submit the DistCp jobs.

To configure the proxy user, configure the Advanced command line options for distcp used in Iceberg Replication = -proxy [***user_name***] key-value pair on the Cloudera Manager Clusters [***Iceberg Replication Service***] Configuration tab.

DMX-3174

Iceberg replication policies fail if the clusters with HDFS HA have different nameservice names and are Auto-TLS enabled on unified realms.

Add the following property for the Advanced configuration snippet for hdfs-site.xml (hdfs_client_config_safety_valve) on the Cloudera Manager Clusters [***HDFS service***] Configuration tab:

mapreduce.job.hdfs-servers.token-renewal.exclude = [***source name service***], [***target name service***].

For more information, see Kerberos setup guidelines.

CDPD-59437

An Iceberg replication policy might not find a table in the database during the replication process if another Iceberg replication policy that is running simultaneously (replicating a different set of tables from the same database) has dropped the table.

OPSAPS-68221: Cloudera Manager Agent installation might fail while upgrading to Cloudera Manager 7.11.3 without installing Python 3 on the Cloudera Manager Server host

Before upgrading to Cloudera Manager 7.11.3, if you do not install Python 3 on the Cloudera Manager Server host, then Cloudera Manager Agent installation might fail. This state is not recoverable by reinstalling Cloudera Manager Agent alone.

- 1. You must uninstall Cloudera Manager Agent package manually.
- Install Python 3 on the host before upgrading to Cloudera Manager 7.11.3. See Installing Python 3.
- **3.** Reinstall the Cloudera Manager Agent.

OPSAPS-68426: Atlas service dependencies are not set during CDH 6 to CDP 7.x.x upgrade if Navigator role instances are not configured under the Cloudera Management Service.

Navigator support has been discontinued in Cloudera Manager 7.11.3. Consequently, if you are using CDH 6 and have Navigator installed, it is necessary to remove the Navigator service before proceeding with the upgrade to Cloudera Manager version 7.11.3 or any higher version. Due to this change, when upgrading the Runtime version from CDH 6 to CDP 7.x.x, it is important to note that Atlas, which replaces Navigator in CDP 7.x.x, might not automatically be set as a service dependency for certain components. The components that could potentially be impacted include: HBase, Hive, Hive on Tez, Hue, Impala, Oozie, Spark, and Sqoop.

Once you have completed the upgrade to CDP 7.x.x and have installed Atlas, it is advised to review and confirm the configuration settings for these services. Specifically, navigate to the respective configuration pages for each service. If you observe that the Atlas dependency is not enabled, you must enable it manually in order to integrate Atlas with that particular service. After adjusting the services' configurations, Cloudera Manager prompts you to restart the services to apply the changes. Note that deploying client configurations might also be necessary as part of this process.

OPSAPS-68340: Zeppelin paragraph execution fails with the User not allowed to impersonate error.

Starting from Cloudera Manager 7.11.3, Cloudera Manager auto-configures the livy_admin_users configuration when Livy is run for the first time. If you add Zeppelin or Knox services later to the existing cluster and do not manually update the service user, the User not allowed to impersonate error is displayed.

If you add Zeppelin or Knox services later to the existing cluster, you must manually add the respective service user to the livy_admin_users configuration in the Livy configuration page.

OPSAPS-68500: The cloudera-manager-installer.bin fails to reach Ubuntu 20 repository on the Archive URL due to redirections.

Agent Installation with Cloudera Manager on Ubuntu20 platform does not function when the selfinstaller method (using the installer.bin file) is employed to install Cloudera Manager. The failure mode is that Cloudera Manager Agent installation step will fail with an error message saying "The repository 'https://archive.cloudera.com/p/cm7/7.11.3/ubuntu2004/apt focal-cm7 InRelease' is not signed."

While adding a cluster in Cloudera Manager and the subsequent agent installation, customers should choose the "Custom Repository" selection, and manually enter the correct repository URL: https:// [credentials]@archive.cloudera.com/p/cm7/7.11.3.0

DMX-3003

The progress.json file is updated along with the progress of the DistCp job run whenever the number of files copied is equal to the incremental count (default is 50) for Iceberg replication policies. The file report does not get synchronized as expected and the reported numbers are also inconsistent.

Click the required Iceberg replication policy on the Cloudera Manager Replication Replication Policies page to see the correct number of files copied for each incremental job run.

DMX-2977, DMX-2978

You cannot view the current status of an ongoing export task (exportCLI) or sync task (syncCLI) for an Iceberg replication policy.

Click the required Iceberg replication policy on the Cloudera Manager Replication Replication Policies page to view the final results of the export task and sync task for the replication policy job run.

OPSAPS-68629: HDFS HTTPFS GateWay is not able to start with custom krb5.conf location set in Cloudera Manager.

On a cluster with a custom krb5.conf file location configured in Cloudera Manager, HDFS HTTPFS role is not able to start because it does not have the custom Kerberos configuration file setting properly propagated to the service, and therefore it fails with a Kerberos related exception: in thread "main" java.io.IOException: Unable to initialize WebAppContext at org.apache.hadoop.http.HttpS erver2.start(HttpServer2.java:1240) at org.apache.hadoop.fs.http.server.HttpFSServerWebServer. start(HttpFSServerWebServer.java:131) at org.apache.hadoop.fs.http.server.HttpFSServerWebS erver.main(HttpFSServerWebServer.java:162) Caused by: java.lang.IllegalArgumentException: Can't get Kerberos realm at org.apache.hadoop.security.HadoopKerberosName.setConfiguration (HadoopKerberosName.java:71) at org.apache.hadoop.security.UserGroupInformation.initialize (UserGroupInformation.java:329) at org.apache.hadoop.security.UserGroupInformation.setConf iguration(UserGroupInformation.java:380) at org.apache.hadoop.lib.service.hadoop.FileSystemA ccessService.init(FileSystemAccessService.java:166) at org.apache.hadoop.lib.server.BaseServic e.init(BaseService.java:71) at org.apache.hadoop.lib.server.Server.initServices(Server.java:581) at org.apache.hadoop.lib.server.Server.init(Server.java:377) at org.apache.hadoop.fs.http.server. HttpFSServerWebApp.init(HttpFSServerWebApp.java:100) at org.apache.hadoop.lib.servlet.Serv erWebApp.contextInitialized(ServerWebApp.java:158) at org.eclipse.jetty.server.handler.Context

Handler.callContextInitialized(ContextHandler.java:1073) at org.eclipse.jetty.servlet.ServletCon textHandler.callContextInitialized(ServletContextHandler.java:572) at org.eclipse.jetty.server.h andler.ContextHandler.contextInitialized(ContextHandler.java:1002) at org.eclipse.jetty.servlet.Se rvletHandler.initialize(ServletHandler.java:765) at org.eclipse.jetty.servlet.ServletContextHandle r.startContext(ServletContextHandler.java:379) at org.eclipse.jetty.webapp.WebAppContext.start Webapp(WebAppContext.java:1449) at org.eclipse.jetty.webapp.WebAppContext.startContext (WebAppContext.java:1414) at org.eclipse.jetty.server.handler.ContextHandler.doStart(ContextHa ndler.java:916) at org.eclipse.jetty.servlet.ServletContextHandler.doStart(ServletContextHandler.j ava:288) at org.eclipse.jetty.webapp.WebAppContext.doStart(WebAppContext.java:524) at or g.eclipse.jetty.util.component.AbstractLifeCycle.start(AbstractLifeCycle.java:73) at org.eclipse.j etty.util.component.ContainerLifeCycle.start(ContainerLifeCycle.java:169) at org.eclipse.jetty.uti l.component.ContainerLifeCycle.doStart(ContainerLifeCycle.java:117) at org.eclipse.jetty.serve r.handler.AbstractHandler.doStart(AbstractHandler.java:97) at org.eclipse.jetty.util.component.Abs tractLifeCycle.start(AbstractLifeCycle.java:73) at org.eclipse.jetty.util.component.ContainerLifeC ycle.start(ContainerLifeCycle.java:169) at org.eclipse.jetty.server.Server.start(Server.java:423) at org.eclipse.jetty.util.component.ContainerLifeCycle.doStart(ContainerLifeCycle.java:110) at org.e clipse.jetty.server.handler.AbstractHandler.doStart(AbstractHandler.java:97) at org.eclipse.jetty. server.Server.doStart(Server.java:387) at org.eclipse.jetty.util.component.AbstractLifeCycle.start (AbstractLifeCycle.java:73) at org.apache.hadoop.http.HttpServer2.start(HttpServer2.java:1218) ... 2 more Caused by: java.lang.IllegalArgumentException: KrbException: Cannot locate default rea lm at java.security.jgss/javax.security.auth.kerberos.KerberosPrincipal.<init>(KerberosPrincipal.j ava:174) at org.apache.hadoop.security.authentication.util.KerberosUtil.getDefaultRealm(Kerberos Util.java:108) at org.apache.hadoop.security.HadoopKerberosName.setConfiguration(HadoopKer berosName.java:69) ...

- 1. Log in to Cloudera Manager.
- **2.** Select the HDFS service.
- 3. Select Configurations tab.
- 4. Search for HttpFS Environment Advanced Configuration Snippet (Safety Valve)
- 5. Add to or extend the HADOOP_OPTS environment variable with the following value: Djava.security.krb5.conf=<the custom krb5.conf location>
- 6. Click Save Changes.

OPSAPS-69897: NPE in Ozone replication from CM 7.7.1 to CM 7.11.3

When you use source Cloudera Manager 7.7.1 and target Cloudera Manager 7.11.3 for Ozone replication policies, the policies fail with Failure during PreOzoneCopyListingCheck execution: null error. This is because the target Cloudera Manager 7.11.3 does not retrieve the required source bucket information for validation from the source Cloudera Manager 7.7.1 during the PreCopyListingCheck command phase. You come across this error when you use source Cloudera Manager versions lower than 7.10.1 and target Cloudera Manager versions higher than or equal to 7.10.1 in an Ozone replication policy.

Upgrade the source Cloudera Manager to 7.11.3 or higher version.

OPSAPS-69481: Some Kafka Connect metrics missing from Cloudera Manager due to conflicting definitions

The metric definitions for kafka_connect_connector_task_metrics_batch_size_avg and kafka_connect_connector_task_metrics_batch_size_max in recent Kafka CSDs conflict with previous definitions in other CSDs. This prevents Cloudera Manager from registering these metrics. It also results in SMM returning an error. The metrics also cannot be monitored in Cloudera Manager chart builder or queried using the Cloudera Manager API.

Contact Cloudera support for a workaround.

OPSAPS-69406: Cannot edit existing HDFS and HBase snapshot policy configuration

The **Edit Configuration** modal window does not appear when you click Actions Edit Configuration on the Cloudera Manager Replication Snapshot Policies page for existing HDFS or HBase snapshot policies. None.

Fixed Issues in Cloudera Manager 7.11.3

Fixed issues in Cloudera Manager 7.11.3.

OPSAPS-65324: The default value of the Cloudera Manager redaction policy configuration CORE_SETTINGS Log and Query Redaction Policy (parameter name: redation_policy) is modified.

The lines Credit Card numbers (with separator) and Social Security numbers (with separator) are modified with the addition of \b symbols before and after the regular expression in the Search field to prevent unintended matching against HDFS block identifiers in the Datanode logs.

Cloudera Manager only applies this change in default value to CDP runtimes with version 7.1.9 and higher.

OPSAPS-47937: Errors while collecting host statistics data, the collectHostStatistics command is frequently timing out

While collecting host statistics data, the collectHostStatistics command is aborting after 150 seconds when the /var/log/messages file is too large. As a result of this, the host statistics data is missing from the diagnostic bundle.

This issue is fixed now by limiting the data size taken from the /var/log/messages file to 300 MB. The collectHostStatistics command now collects only 300 MB of the latest data from the /var/log/messages file to avoid timeouts.

OPSAPS-68044: Certain Cloudera Runtime services (such as HDFS) might fail to start on RHEL 8.8 with FIPS mode enabled.

While configuring Cloudera Manager cluster for installation or upgrade process on RHEL 8.8 with FIPS mode enabled, certain Cloudera Runtime services such as HDFS might fail to start and throws the following error:OpenSSL internal error: FATAL FIPS SELFTEST FAILURE

This issue is fixed now.

OPSAPS-66052: Cloudera Manager is unable to execute certain operations when you enable the noexec option for the /tmp directory

When you enable noexec option for the /tmp directory of the cluster hosts, Cloudera Manager is not able to complete some operations, most notably for the Add Hosts workflow and while generating TLS certificates.

This behavior is resolved now and Cloudera Manager functions normally when you enable the noex ec option for the /tmp directory on cluster hosts.

OPSAPS-67942: Installation failed due to schematool error

Setting the hive.hook.proto.base-directory for Hive Metastore (HMS) in hive-site.xml is causing sys.db creation to fail because of incompatibility issues between Cloudera Manager 7.11.3 and CDH 7.1.7 SP1/SP2. This patch addresses the issue and sets the above configuration only if the CDH version of Hive is atleast CDH 7.1.8.

OPSAPS-67968: An issue while upgrading to Cloudera Manager 7.11.3 without upgrading the Cloudera Runtime version 7.1.7 SP2

With this fix, you can now upgrade to Python 3 compliant Cloudera Manager 7.11.3 without upgrading the Cloudera Runtime version (Cloudera Runtime 7.1.7 SP2 and below versions - which are not Python 3 compliant). Cloudera Manager 7.11.3 now supports Cloudera Runtime 7.1.7 SP2 and below versions.



Note: Cloudera Runtime 7.1.7 SP2 and below versions are only supported on Python 2.7.

Cloudera Manager 7.11.3 typically supports Python 3.8 version. Cloudera Manager 7.11.3 also supports Python 3.9 version when running on RHEL 9.1 operating system.

OPSAPS-63724

By default, the snapshot diff-based (incremental) HDFS - HDFS replication falls back to bootstrap (full file listing / FFL) replication when there are unexpected target-side changes. By enabling this workaround, certain target-side changes are tolerated by incremental replication without falling back to FFL. Note that when source side HDFS moves are expected to be synchronized the workaround mentioned in OPSAPS-66197 is recommended to be activated.

Activating this workaround:

- Set "com.cloudera.enterprise.distcp.check-for-safe-to-merge-target-side-changes.enabled" to "true" in the "YARN Service Advanced Configuration Snippet (Safety Valve) for core-site.xml" on the destination side, and then restart the stale services / redeploy client configuration. (Note that enabling OPSAPS-66197 uses a different advanced configuration snippet).
- In an incremental replication run, check the stderr log of the first "Run Pre-Filelisting Check" and make sure the INFO distcp.PreCopyListingCheck: Check for safe to ignore (merge) target side changes is enabled. message appears.

Usage notes:

• When a safe-to-ignore target change is found, "Run Pre-Filelisting Check" prints the following messages to its stderr log:

```
INFO util.DistCpUtils: There are changes on target, falling
back to regular distcp
NFO distcp.PreCopyListingCheck: The changes on target are safe
to ignore.
INFO distcp.PreCopyListingCheck: Note that it is up to the
downstream processing steps if it falls back to full file li
sting or continue with snapshot diff execution
INFO distcp.PreCopyListingCheck: Changes to target: true
INFO distcp.PreCopyListingCheck: Changes to target are safe to
ignore: true
```

• When target changes are not found safe-to-ignore, then the details about the reason appears in the messages:

```
INFO distcp.PreCopyListingCheck: Changes to target: true
INFO distcp.PreCopyListingCheck: Changes to target are safe
to ignore: false
```

Allowed changes:

The following destination side changes (snapshot diff entries) are considered safe-to-ignore when this workaround is enabled:

- Additions (+): only if they are empty directories or contain only directories, all present on the source as directory.
- Deletions (): only the source side path also missing.
- Modifications (M): must have an immediate, allowed (+) or (-) child path.

OPSAPS-63930

By default, snapshot diff-based (incremental) HDFS - HDFS replication uses a temp directory, created in the parent of replication destination directory to synchronize source-side rename and delete operations: deleted and renamed paths are first moved into this temporary directory, then the renamed ones will be moved to their target followed by the deletion of this temporary directory (thus deleting the paths scheduled to be deleted). Note that OPSAPS-63759 provides an optional behavior to execute individual deletes without these moves.

This behavior of incremental replication leads to failure and fallback to bootstrap (full file listing) replication when the replication process can not create this temporary directory (due to restrictive HDFS permissions) or when the replication destination contains one or more HDFS encryption zones (because HDFS moves can not cross encryption zone boundary).

This optional workaround solves these problems by executing rename operations in-place when possible, otherwise using the best possible temporary rename operations without the need of the above mentioned common temporary directory. Note that this workaround can be considered as a superset of OPSAPS-63759. That is when both are enabled, the current one is applied.

Activating this workaround:

- Set HDFS service core-site.xml advanced configuration snippet (on the destination side) "com.cloudera.enterprise.distcp.direct-rename-and-delete.enabled" to "true".
- In an incremental replication run, check the stderr log of the last "Trigger a HDFS replication job on one of the available HDFS roles." step, and make sure the INFO distcp.DistCpSync: Will use direct rename and delete (for non cloud target) when using snapshot diff based sync. Te mp directory creation on the target will be skipped. message is displayed.

Adjusting delete logging: By default, every 100000 direct delete operations executed by this workaround are logged. This is useful for following the synchronization of large source side deletes. This default interval can be overridden by setting the "com.cloudera.enterprise.distcp.direct-delete.log-interval" advanced configuration snippet to an integer value greater than 0. Note that this advanced configuration snippet is shared with a workaround in OPSAPS-63759.

Usage notes: There can be conflicting source side renames and rename - delete interactions when their destination side replay need to use temporary renames (for example, a name swap between two paths using three renames). For these cases, the temporary rename destination will typically be next to the final rename destination (will share the same parent path) avoiding both above mentioned failure scenarios. Such temporary renames will be logged during execution like:

```
distcp.DistCpSync: Executing a temp rename: /test-repl-target/te
st-repl-source/file2 -> /test-repl-target/test-repl-source/file2
748016654
```

After execution, the number of operations will also be logged like:

```
INFO distcp.DistCpSync: Synced 0 through-tmp/cloud rename(s) and
0 through-tmp delete(s) to target.
INFO distcp.DistCpSync: Synced 2 direct delete(s) to target.
INFO distcp.DistCpSync: Synced 2 direct rename(s) to target.
INFO distcp.DistCpSync: Used 2 additional temporary rename(s)
during syncing.
```

OPSAPS-66197

Snapshot diff-based (incremental) HDFS to HDFS replication might corrupt destination directory structure when:

- there is a source side HDFS move/rename operation.
- the (move/rename) target on the replication destination is an existing unexpected directory.

OPSAPS-63724 introduced an optional workaround where the target-side directory creations are ignored. When a colliding source-side move is expected both workarounds are recommended to be activated.

Workaround:

• Set the HDFS service core-site.xml advanced configuration snippet (also called safety valve) (on the destination side) "com.cloudera.enterprise.distcp.overwrite-merge-existing-rename-targets.enabled" to "true". (Note that enabling the workaround in OPSAPS-63724 uses a different advanced configuration snippet).

• In an incremental replication run, check the stderr log of the last "Trigger a HDFS replication job on one of the available HDFS roles." step and make sure that the INFO distcp.DistCpSync: Overwrite merge of already existing move targets is enabled message is displayed.

Usage notes:

- When there is a conflicting replicated source side move/rename operation where on the destination side the target exists, there will be a merge attempt:
- When the source side moved path is a directory and the conflicting destination side path is also a directory their contents will be merged.
- When the destination side conflicting path is a file it will be overwritten by the replicated move.
- When the source side moved path is a file the destination side conflicting path will be overwritten by the replicated move.
- In case of other failures replication is expected to fall back to bootstrap (full file listing) run.

Details of merge activity (when there is a conflicting path) is logged in the same stderr log with messages containing INFO distcp.DistCpSync\$OverwriteMergeRenameBehavior.

OPSAPS-63558

Previously, DistCp did not correctly report renames and deletes in case of snapshot diff-based HDFS replications. This change extends DistCp's output report to contain counters related to snapshot diff-based replications beside the already reported counters. These counters are added to the following group: com.cloudera.enterprise.distcp.DistCpSyncCounter.

The following new counters are added:

- FILES_MOVED_TO_COMMON_TEMP_DIR: Number of files and directories moved to a common temporary directory to be renamed or deleted later in the process. This counter is the sum of FILES_DELETED_VIA_COMMON_TEMP_DIR and FILES_RENAMED_VIA_COMMON_TEMP_DIR.
- FILES_DELETED_VIA_COMMON_TEMP_DIR: Number of files moved to a common temporary directory to be deleted later.
- FILES_RENAMED_VIA_COMMON_TEMP_DIR: Number of files moved to a common temporary directory first, then moved to their final place.
- FILES_DIRECT_DELETED: Number of files deleted directly. This is a feature introduced in OPSAPS-63759.
- FILES_DIRECT_RENAMED: Number of files renamed directly, without moving to an intermediate temporary directory. This is a feature introduced in OPSAPS-63930.
- FILES_DIRECT_RENAMED_VIA_TEMP_LOCATION: Number of files moved to an intermediate temporary directory and then renamed. This intermediate temporary directory is different from the common temporary directory referenced in the FILES_RENAMED_VIA_COMMON_TEMP_DIR counter's description. This is also related to OPSAPS-63930.

The common temporary directory is a sibling of the replication target directory.

The values of FILES_DELETED_VIA_COMMON_TEMP_DIR and FILES_DIRECT_DELETED are also aggregated in the replication result as the number of files deleted.

OPSAPS-65831: DistCp job deletes multiple threads for bootstrap replication

Performance of bootstrap or FFL (full file listing) replication for destination-side delete of paths missing from the source is improved with the following optional behaviors.

• FFL replication schedules all the missing paths for deletion regardless of parent relationships. When the com.cloudera.enterprise.distcp.parent-only-delete.enabled safety valve is set to "true", only the topmost deleted paths are scheduled for deletions and their descendants or children cannot be accessed. This is optional and by default turned off (which preserves the previous behavior).

- Delete requests can be issued from multiple threads concurrently to improve performance, and can be enabled and configured using the following safety valves:
 - com.cloudera.enterprise.distcp.parallel-ffl-delete.enabled. Default is "false".
 - com.cloudera.enterprise.distcp.parallel-ffl-delete.threads. Default is 20.
 - com.cloudera.enterprise.distcp.parallel-ffl-delete.max-queuesize.Default is 10000.

OPSAPS-65823

Added periodic progress logging during copy listing. Optionally, the performance statistics of file system operations (min/max/avg/total time since last log and since beginning of copy listing) are also printed.

When the bootstrap (full file listing) run launches target side copy listing (to handle deletions) the reducer log also contains the log messages of the reducer activity. Overview of this activity (status reducer step; listed path count) is also logged on the main DistCp process. Job counters containing reducer timing measurements and listed target side path count are also added.

By default, the interval of copy listing logging is 10 seconds which can be adjusted by setting the com.cloudera.enterprise.distcp.copy-listing.progress-log.interval.seconds configuration parameter in the HDFS replication core-site.xml configuration.

Setting detailed log is done by setting the com.cloudera.enterprise.distcp.copylisting.detailed-progress-log.enabled configuration parameter to "true".

Disabling progress logging is done by setting the com.cloudera.enterprise.distcp.copy-listing.basic-progress-log.enabled configuration parameter to "false".

For testing purposes, the poll interval to check the progress of the MR job (from DistCp main process) can be set with the com.cloudera.enterprise.distcp.job-poll-interval.seconds configuration parameter.

OPSAPS-65104

Importing table column statistics for Hive replication is thread-safe but causes performance regression.

To resolve this issue, perform the following steps:

- 1. Go to the Cloudera Manager Clusters [*** Hive Service ***] Configuration tab.
- 2. Locate the hive_replication_env_safety_valve property.
- 3. Add only one of the following key-value pair depending on your requirement:
 - COLUMN_STATS_IMPORT_MULTI_THREADED=true

This ensures that the column statistics import operation is multi-threaded for Hive replication.

• SKIP_COLUMN_STATS_IMPORT=true

This ensures that the column statistics import is skipped entirely.

OPSAPS-66517: Changing password from Home username Change Password bypasses validation

In Cloudera Manager, while changing the password for the current user from Home username Change Password, password validations are completely bypassed. This issue is fixed now and it now validates the password before saving the new password.

OPSAPS-67490: Cloudera Manager unable to deploy the Hadoop User Group Mapping LDAP Bind User Password configuration completely

Fixed an issue where Cloudera Manager is unable to deploy complete configurations from Core Configurations (CORE_SETTINGS-1) to client configurations under local /etc directory in the JCEKS file.

OPSAPS-65267

Cross-site sessions were prohibited in the latest browsers because of SameSite header by default was set to Lax. This issue is fixed now by adding SameSite=None with a secure attribute for the session cookies that are created after login so that cross-site secure cookies are supported.

The secure attribute works only with TLS-configured clusters. You must have a TLS-enabled cluster for cross-site sessions to work.

OPSAPS-66080: Optimize pattern.compile in CspUtils.java

When Cloudera Manager is running, compiling the regex pattern for CSP multiple times causes other threads to wait, and that results in the slowness of Cloudera Manager. This issue is fixed now.

OPSAPS-66198: On Cloudera Manager UI, the Install Oozie ShareLib command fails to install shared libraries for the Oozie service

On Cloudera Manager UI, the Install Oozie ShareLib command fails to install shared libraries for the Oozie service if you configure the Kerberos krb_krb5_conf_path file path at the non-default file path. This issue is fixed now.

OPSAPS-67152: Cloudera Manager does not allow you to update some configuration parameters.

Cloudera Manager does not allow you to set to "0" for the dfs_access_time_precision and dfs_name node_accesstime_precision configuration parameters.

You will not be able to update dfs_access_time_precision and dfs_namenode_accesstime_precision to "0". If you try to enter "0" in these configuration input fields, then the field gets cleared off and results in a validation error: This field is required. This issue is fixed now.

OPSAPS-66023: Error message about an unsupported ciphersuite while upgrading or installing cluster with the latest FIPS compliance

When upgrading or installing a FIPS enabled cluster, Cloudera Manager is unable to download the new CDP parcel from the Cloudera parcel archive.

Cloudera Manager displays the following error message:

HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA

This issue is fixed now by correcting the incorrect ciphersuite selection.

OPSAPS-67897, OPSAPS-68023

Ozone replication policies do not fail and the files on the target cluster are deleted successfully when you set the Advanced Options Delete Policy option to 'Delete to Trash' or 'Delete Permanently' during the Ozone replication policy creation process in CDP Private Cloud Base Replication Manager UI, or if you set the "removeMissingFiles" parameter to 'true' while creating the Ozone replication policy using Cloudera Manager REST APIs.

Fixed Common Vulnerabilities and Exposures in Cloudera Manager 7.11.3 and Cloudera Manager 7.11.3 cumulative hotfixes

Common Vulnerabilities and Exposures (CVE) that are fixed in Cloudera Manager 7.11.3 and Cloudera Manager 7.11.3 cumulative hotfixes.

Cloudera Manager 7.11.3 CHF5

- CVE-2020-11971 Apache Camel
- CVE-2023-43642 Snappy-java
- CVE-2022-22965 Spring Framework
- CVE-2023-20860 Spring Framework
- CVE-2022-22950 Spring Framework
- CVE-2022-22971 Spring Framework
- CVE-2023-20861 Spring Framework
- CVE-2023-20863 Spring Framework
- CVE-2022-22968 Spring Framework
- CVE-2022-22970 Spring Framework
- CVE-2021-22060 Spring Framework
- CVE-2021-22096 Spring Framework

Cloudera Manager 7.11.3 CHF4

No Common Vulnerabilities and Exposures (CVE) are fixed in Cloudera Manager 7.11.3 CHF4.

Cloudera Manager 7.11.3 CHF3

- CVE-2022-46364 Apache CXF
- CVE-2022-46363 Apache CXF
- CVE-2022-1415 Drools
- CVE-2021-41411 Drools
- CVE-2022-41853 Hsqldb
- CVE-2011-4461 Mortbay Jetty
- CVE-2009-1523 Mortbay Jetty
- CVE-2009-4611 Mortbay Jetty
- CVE-2009-5048 Mortbay Jetty
- CVE-2009-5049 Mortbay Jetty
- CVE-2009-4609 Mortbay Jetty
- CVE-2009-1524 Mortbay Jetty
- CVE-2009-4610 Mortbay Jetty
- CVE-2009-4612 Mortbay Jetty
- CVE-2021-35515 Commons-Compress
- CVE-2021-35516 Commons-Compress
- CVE-2021-35517 Commons-Compress
- CVE-2021-36090 Commons-Compress
- CVE-2023-25613 Apache Kerby
- CVE-2022-41915 Netty
- CVE-2018-11799 Apache Oozie
- CVE-2017-15712 Apache Oozie
- CVE-2022-34169 Xalan
- CVE-2023-34453 Snappy-java
- CVE-2023-34454 Snappy-java
- CVE-2023-34455 Snappy-java
- CVE-2023-34034 Spring Security
- CVE-2020-13936 Apache Velocity

Cloudera Manager 7.11.3 CHF2

• CVE-2022-25647 - Gson

- CVE-2021-28165 Eclipse Jetty
- CVE-2022-2048 Eclipse Jetty
- CVE-2020-27223 Eclipse Jetty
- CVE-2021-28169 Eclipse Jetty
- CVE-2021-34428 Eclipse Jetty
- CVE-2021-28163 Eclipse Jetty
- CVE-2022-2047 Eclipse Jetty
- CVE-2022-45688 org.json
- CVE-2023-5072 org.json
- CVE-2023-3635 Okio

Cloudera Manager 7.11.3 CHF1

No Common Vulnerabilities and Exposures (CVE) are fixed in Cloudera Manager 7.11.3 CHF1.

Cloudera Manager 7.11.3

- CVE-2021-25642
- CVE-2022-25168
- CVE-2022-31129
- CVE-2021-36373
- CVE-2021-36374
- CVE-2020-9493
- CVE-2022-23305
- CVE-2023-26464
- CVE-2018-14721
- CVE-2018-14718
- CVE-2018-14719
- CVE-2018-14720
- CVE-2018-19360
- CVE-2018-19361
- CVE-2018-19362
- CVE-2018-12022
- CVE-2018-12023
- CVE-2022-36364
- CVE-2017-15095
- CVE-2018-5968
- CVE-2020-28491
- CVE-2022-40146
- CVE-2022-41704
- CVE-2022-42890
- CVE-2022-38398
- CVE-2022-38648
- CVE-2020-15522
- CVE-2020-0187
- CVE-2020-26939
- CVE-2020-13955
- CVE-2019-14862
- CVE-2023-24998
- CVE-2022-23457
- CVE-2022-24891

- CVE-2018-11792
- CVE-2021-28131
- CVE-2018-11785
- CVE-2022-21724
- CVE-2022-26520
- CVE-2022-31197
- CVE-2022-41946
- CVE-2021-27905
- CVE-2021-44548
- CVE-2021-29943
- CVE-2020-13941
- CVE-2017-3163
- CVE-2017-3164
- CVE-2018-1308
- CVE-2019-12401
- CVE-2019-0193
- CVE-2015-8795
- CVE-2015-8796
- CVE-2015-8797
- CVE-2018-11802
- CVE-2020-5421
- CVE-2022-22978
- CVE-2021-22112
- CVE-2022-22976
- CVE-2016-1000027
- CVE-2020-5397
- CVE-2022-40152
- CVE-2022-40151
- CVE-2022-41966
- CVE-2020-10683
- CVE-2018-1000632
- CVE-2014-0229
- CVE-2014-3627
- CVE-2014-3566
- CVE-2013-4221
- CVE-2013-4271
- CVE-2017-14868
- CVE-2017-14949
- CVE-2014-1868
- CVE-2018-8010
- CVE-2018-8026
- CVE-2017-5637
- CVE-2021-37533
- CVE-2018-14040
- CVE-2022-37865
- CVE-2022-37866
- CVE-2013-2035
- CVE-2014-125087
- CVE-2021-33813
- CVE-2014-3643

- CVE-2022-40149
- CVE-2022-40150
- CVE-2022-45685
- CVE-2022-45693
- CVE-2023-1436
- CVE-2017-7657
- CVE-2017-7658
- CVE-2017-7656
- CVE-2017-9735
- CVE-2020-27216
- CVE-2019-10247
- CVE-2019-10241
- CVE-2018-12536
- CVE-2019-10246
- CVE-2016-5725
- CVE-2023-1370
- CVE-2021-37714
- CVE-2022-36033
- CVE-2016-4000
- CVE-2018-1320
- CVE-2019-0205
- CVE-2019-0210
- CVE-2018-11798
- CVE-2019-17571
- CVE-2021-4104
- CVE-2016-2402
- CVE-2021-27807
- CVE-2021-27906
- CVE-2021-31811
- CVE-2021-31812
- CVE-2022-26336
- CVE-2022-1471
- CVE-2017-18640
- CVE-2022-25857
- CVE-2022-38749
- CVE-2022-38751
- CVE-2022-38752
- CVE-2022-41854
- CVE-2022-38750
- CVE-2017-8028
- CVE-2019-11272
- CVE-2019-3795
- CVE-2019-14887
- CVE-2022-23437
- CVE-2020-11988

Deprecation notices in Cloudera Manager 7.11.3

Certain features and functionalities have been removed or deprecated in Cloudera Manager 7.11.3. You must review these items to understand whether you must modify your existing configuration. You can also learn about the features that will be removed or deprecated in the future release to plan for the required changes.

Terminology

Items in this section are designated as follows:

Deprecated

Technology that Cloudera is removing in a future Cloudera Manager release. Marking an item as deprecated gives you time to plan for removal in a future Cloudera Manager release.

Moving

Technology that Cloudera is moving from a future Cloudera Manager release and is making available through an alternative Cloudera offering or subscription. Marking an item as moving gives you time to plan for removal in a future Cloudera Manager release and plan for the alternative Cloudera offering or subscription for the technology.

Removed

Technology that Cloudera has removed from Cloudera Manager and is no longer available or supported as of this release. Take note of technology marked as removed since it can potentially affect your upgrade plans.

Deprecation Notices for Cloudera Manager

Certain features and functionality are deprecated or removed in Cloudera Manager 7.11.3. You must review these changes along with the information about the features in Cloudera Manager that will be removed or deprecated in a future release.

Removed

SHA-1-GNU Privacy Guard (GPG) keys

The SHA-1 hashing algorithm based GPG signing keys are removed. Instead, use a stronger, secure hashing algorithm called SHA-256.

Platform and OS

The listed Operating Systems and databases are deprecated or removed from the Cloudera Manager 7.11.3 release.

Database Support:

The listed databases are deprecated from the 7.11.3 release.

- Postgres 10 (FIPS)
- MariaDB 10.3
- MariaDB 10.2
- Oracle 12
- MySQL 5.6

Operating System

The listed operating system is removed from the 7.11.3 release.

• Ubuntu 18.04



Note: Ubuntu 18 Operating System is not supported from Cloudera Manager 7.11.3 to Cloudera Manager 7.11.3 CHF4 versions. You must upgrade the Operating System from Ubuntu 18 to Ubuntu 20 before you upgrade to Cloudera Manager 7.11.3 CHF4. For performing major OS upgrade, see Upgrading the Operating System to a new Major Version.

Cumulative hotfixes

You can review the list of cumulative hotfixes that were shipped for Cloudera Manager 7.11.3 release.

Cloudera Manager 7.11.3 Cumulative hotfix 5

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 5.

This cumulative hotfix was released on April 8, 2024.



Note: Contact Cloudera Support for questions related to any specific hotfixes.

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF5 (version: 7.11.3.7-52024171):

OPSAPS-68340: Zeppelin paragraph execution fails with the User not allowed to impersonate error.

Starting from Cloudera Manager 7.11.3, Cloudera Manager auto-configures the livy_admin_users configuration when Livy is run for the first time. If you add Zeppelin or Knox services later to the existing cluster and do not manually update the service user, the User not allowed to impersonate error is displayed.

If you add Zeppelin or Knox services later to the existing cluster, you must manually add the respective service user to the livy_admin_users configuration in the Livy configuration page.

OPSAPS-69847:Replication policies might fail if source and target use different Kerberos encryption types

Replication policies might fail if the source and target Cloudera Manager instances use different encryption types in Kerberos because of different Java versions. For example, the Java 11 and higher versions might use the *aes256-cts* encryption type, and the versions lower than Java 11 might use the *rc4-hmac* encryption type.

Ensure that both the instances use the same Java version. If it is not possible to have the same Java versions on both the instances, ensure that they use the same encryption type for Kerberos. To check the encryption type in Cloudera Manager, search for krb_enc_types on the Cloudera Manager Administration Settings page.

Following are the list of fixed issues that were shipped for Cloudera Manager 7.11.3 CHF5 (version: 7.11.3.7-52024171):

OPSAPS-70077: Need the ECS Update Ingress Controller Certificate action added to the API

The ECS Update Ingress Controller Certificate action has been added to the API.

OPSAPS-70084: Upgrade ingress controller cert command failed for DSA encrypted private key

DSA has been dropped as a supported key type for ingress certificate private keys.

OPSAPS-69668: Support password protected ingress private key

Passwords are now supported for ingress certificate private keys.

OPSAPS-69808: Update AuthzMigrator GBN to point to latest non-expired GBN

Users will now be able to export sentry data only for given Hive objects (databases, tables, and the respective URLs) by using the config authorization.migration.export.migration_objects during export.

OPSAPS-69057: Customizable authorization-migration-site.xml for Sentry-to-Ranger migration

You can now add additional arguments to override any existing property in the authorizationmigration-site.xml file. The Sentry to Ranger migration process during the Hive replication policy run uses this file. These additional arguments are used during the Sentry to Ranger migration process for Sentry export on the source and Ranger import on the destination. You can enter the arguments using the CM API body as shown in the following sample snippet:

```
"hiveArguments": {
    ...
    "rangerImportProperties": {
        "authorization.migration.destination.location.prefix": "
hdfs://nameservice",
        "some.other.prop": "some_property"
    },
        "sentryExportProperties": {
            "authorization.migration.role.permissions": "true",
            "export.prop": "export_prop_sentry",
            "authorization.migration.destination.location.prefix":
"hdfs://nameservice"
        },
        ...
}
```

OPSAPS-69207: Customizable authorization-migration-site.xml for Sentry-to-Ranger migration

During the Hive external table replication creation process, you can modify the properties in the authorization-migration-site.xml file on the **Sentry-Ranger Migration** tab. This tab appears after you choose the If Sentry permissions were exported from the CDH cluster, import both Hive object and URL permissions or If Sentry permissions were exported from the CDH cluster, import only Hive object permissions option in the Hive external table replication policy wizard General Permissions field.

OPSAPS-69709: Set Sqoop Atlas hook to send notifications synchronously

Sqoop has an Atlas hook which by default runs asynchronously to send notifications to the Atlas server. In certain cases, the Java Virtual Machine (JVM) in which Sqoop is running can shut down before the Kafka notification of the Atlas hook is sent. This can result in lost notifications.

This issue is fixed by ensuring that the notifications are synchronous.

OPSAPS-69759: Multiple TestDFSIO(Mapreduce job) failure during COD ZDU

This issue has been fixed and Mapreduce job failures will no longer occur.

OPSAPS-69846: Ozone multitenancy PutObject throws Internal Server Error with linked and encrypted bucket

If Ozone is installed with custom Kerberos principals for its roles, operations on encrypted buckets can fail as Ranger KMS does not have its proxy users and groups configured for the custom s3 gateway user.

This issue is fixed now. From 7.11.3 CHF5 onwards, you do not need to manually configure the s3g proxy user for KMS.

The repositories for Cloudera Manager 7.11.3-CHF5 are listed in the following table:

Table 1: Cloudera Manager 7.11.3-CHF5

Repository Type	Repository Location
RHEL 9 Compatible	Repository:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.7/redhat9/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.7/redhat9/yum/cloudera-manager.repo
RHEL 8 Compatible	Repository:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.7/redhat8/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.7/redhat8/yum/cloudera-manager.repo
RHEL 7 Compatible	Repository:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.7/redhat7/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.7/redhat7/yum/cloudera-manager.repo
SLES 15	Repository:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.7/sles15/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.7/sles15/yum/cloudera-manager.repo
SLES 12	Repository:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.7/sles12/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.7/sles12/yum/cloudera-manager.repo
	<pre>https://username:password@archive.cloudera.com/p/ cm7/7.11.3.7/sles12/yum Repository File: https://username:password@archive.cloudera.com/p/ cm7/7.11.3.7/sles12/yum/cloudera-manager.repo</pre>

Repository Type	Repository Location
Ubuntu 20	Repository:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.7/ubuntu2004/apt
	Repository File:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.7/ubuntu2004/apt/cloudera-manager.list
IBM PowerPC RHEL 7	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.7/redhat7-ppc/yum
IBM PowerPC RHEL 8	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.7/redhat8-ppc/yum

Cloudera Manager 7.11.3 Cumulative hotfix 4

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 4.

This cumulative hotfix was released on March 8, 2024.



Note: Contact Cloudera Support for questions related to any specific hotfixes.

New features and changed behavior for Cloudera Manager 7.11.3 CHF4 (version: 7.11.3.6-50817646): FIPS support for JDK11 in Zeppelin

Added FIPS support for JDK11 in Zeppelin.

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF4 (version: 7.11.3.6-50817646):

OPSAPS-68340: Zeppelin paragraph execution fails with the User not allowed to impersonate error.

Starting from Cloudera Manager 7.11.3, Cloudera Manager auto-configures the livy_admin_users configuration when Livy is run for the first time. If you add Zeppelin or Knox services later to the existing cluster and do not manually update the service user, the User not allowed to impersonate error is displayed.

If you add Zeppelin or Knox services later to the existing cluster, you must manually add the respective service user to the livy_admin_users configuration in the Livy configuration page.

OPSAPS-69342: Access issues identified in MariaDB 10.6 were causing discrepancies in High Availability (HA) mode

MariaDB 10.6, by default, includes the property require_secure_transport=ON in the configuration file (/etc/my.cnf), which is absent in MariaDB 10.4. This setting prohibits non-TLS connections, leading to access issues. This problem is observed in High Availability (HA) mode, where certain operations may not be using the same connection.

To resolve the issue temporarily, you can either comment out or disable the line require_secure_t ransport in the configuration file located at /etc/my.cnf.

OPSAPS-68452: Azul Open JDK 8 and 11 are not supported with Cloudera Manager

Azul Open JDK 8 and 11 are not supported with Cloudera Manager. To use Azul Open JDK 8 or 11 for Cloudera Manager RPM/DEBs, you must manually create a symlink between the Zulu JDK installation path and the default JDK path.

After installing Azul Open JDK8 or 11, you must run the following commands on all the hosts in the cluster:

Azul Open JDK 8

RHEL or SLES

```
# sudo ln -s /usr/lib/jvm/java-8-zulu-openjdk-jdk /usr/lib/jvm/j
ava-8-openjdk
```

Ubuntu or Debian

```
# sudo ln -s /usr/lib/jvm/zulu-8-amd64 /usr/lib/jvm/java-8-openj
dk
```

Azul Open JDK 11

For DEBs only

sudo ln -s /usr/lib/jvm/zulu-11-amd64 /usr/lib/jvm/java-11

CDPD-62834: Status of the deleted table is seen as ACTIVE in Atlas after the completion of navigator2atlas migration process

The status of the deleted table displays as ACTIVE.

None

CDPD-62837: During the navigator2atlas process, the hive_storagedesc is incomplete in Atlas

For the hive_storagedesc entity, some of the attributes are not getting populated.

None

OPSAPS-69897: NPE in Ozone replication from CM 7.7.1 to CM 7.11.3

When you use source Cloudera Manager 7.7.1 and target Cloudera Manager 7.11.3 for Ozone replication policies, the policies fail with Failure during PreOzoneCopyListingCheck execution: null error. This is because the target Cloudera Manager 7.11.3 does not retrieve the required source bucket information for validation from the source Cloudera Manager 7.7.1 during the PreCopyListingCheck command phase. You come across this error when you use source Cloudera Manager versions lower than 7.10.1 and target Cloudera Manager versions higher than or equal to 7.10.1 in an Ozone replication policy.

Upgrade the source Cloudera Manager to 7.11.3 or higher version.

Following are the list of fixed issues that were shipped for Cloudera Manager 7.11.3 CHF4 (version: 7.11.3.6-50817646):

OPSAPS-69387: Update Spark 3 parcel CSD's repository URL to point to CDP 7.1.9.x cluster in CM

Updated the Spark 3 parcel's repository URL to point to https://archive.cloudera.com/p/spark3/3.3.7190.0/parcels/ instead of https://archive.cloudera.com/p/spark3/3.3.7180.0/parcels/.

OPSAPS-69711: Cloudera Manager - ECS Server host on RHEL 9.1 keeps getting entropy alerts

The host level entropy health test turns into BAD state if the OS version is among RHEL, Cent OS. OEL 9.x. That can cause BAD health state for the services deployed into these hosts. This issue is fixed now.

OPSAPS-69458: Custom properties atlas.jaas.KafkaClient.option.password appears in a clear text in CDP cluster services.

CDP Private Cloud Base 7.1.9 cluster had a configuration property with a clear text password which is a Information security breach. The password is now masked or encrypted in the cluster.

OPSAPS-69480: Hardcode MR add-opens-as-default config

Cloudera Manager uses fixed runtime versions when determining clients, instead of using the one connected to the deployed runtime version, which can cause issues. During an upgrade if an app is submitted with a client containing MAPREDUCE-7449 to a runtime that doesn't contain MAPREDUCE-7449's related changes, the application submission fails. To fix this issue MAPREDUCE-7468 changes the default behaviour of the feature to avoid including the placeholder by default. Cloudera Manager has a hardcoded property from the runtime versions where the replacement is correctly done in NM code.

OPSAPS-69481: Some Kafka connect metrics missing from Cloudera Manager due to conflicting definitions

Cloudera Manager now registers kafka_connect_connector_task_metrics_batch_size_avg and kafk a_connect_connector_task_metrics_batch_size_max metrics correctly.

OPSAPS-69556: While upgrading from CDP Private Cloud Data Services 1.5.1 to 1.5.2, the public registry with public bits fails with ImagePull Errors, and the docker registry modified to point to docker-private during the upgrade

Previously, when upgrading using the Cloudera public registry with public bits, the Docker registry would incorrectly change to point to docker-private.infra.cloudera.com. This issue is now fixed to point to the correct registry.

OPSAPS-69357: Yarn application bundle script needs to be backwards compatible with python 2.7.

Application bundle collection has been fixed to support both Python2 and Python3 environments.

OPSAPS-68288: Cloudera Manager waits on "Refreshing Resource manager" during the time when the node-manager is being decommissioned

The decommission now works as expected.

OPSAPS-69502: Upgrade failures from CDH6 to 7.1.7 SP3 because ACL is not the expected for znode

Updated the zk-client.sh to follow the output change of the ZK CLI during upgrade so that the upgrade no longer fails.

The repositories for Cloudera Manager 7.11.3-CHF4 are listed in the following table:

Repository Type	Repository Location
RHEL 9 Compatible	Repository:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.6/redhat9/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.6/redhat9/yum/cloudera-manager.repo
RHEL 8 Compatible	Repository:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.6/redhat8/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.6/redhat8/yum/cloudera-manager.repo

Table 2: Cloudera Manager 7.11.3-CHF4

Repository Type	Repository Location
RHEL 7 Compatible	Repository:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.6/redhat7/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.6/redhat7/yum/cloudera-manager.repo
SLES 15	Repository:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.6/sles15/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.6/sles15/yum/cloudera-manager.repo
SLES 12	Repository:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.6/sles12/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.6/sles12/yum/cloudera-manager.repo
Ubuntu 20	Repository:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.6/ubuntu2004/apt
	Repository File:
	https://username:password@archive.cloudera.com/p/ cm7/7.11.3.6/ubuntu2004/apt/cloudera-manager.list

Cloudera Manager 7.11.3 Cumulative hotfix 3

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 3.

This cumulative hotfix was released on February 23, 2024.



Note: Contact Cloudera Support for questions related to any specific hotfixes.

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF3 (version: 7.11.3.4-50275000):

OPSAPS-68340: Zeppelin paragraph execution fails with the User not allowed to impersonate error.

Starting from Cloudera Manager 7.11.3, Cloudera Manager auto-configures the livy_admin_users configuration when Livy is run for the first time. If you add Zeppelin or Knox services later to the

existing cluster and do not manually update the service user, the User not allowed to impersonate error is displayed.

If you add Zeppelin or Knox services later to the existing cluster, you must manually add the respective service user to the livy_admin_users configuration in the Livy configuration page.

OPSAPS-69342: Access issues identified in MariaDB 10.6 were causing discrepancies in High Availability (HA) mode

MariaDB 10.6, by default, includes the property require_secure_transport=ON in the configuration file (/etc/my.cnf), which is absent in MariaDB 10.4. This setting prohibits non-TLS connections, leading to access issues. This problem is observed in High Availability (HA) mode, where certain operations may not be using the same connection.

To resolve the issue temporarily, you can either comment out or disable the line require_secure_t ransport in the configuration file located at /etc/my.cnf.

OPSAPS-68452: Azul Open JDK 8 and 11 are not supported with Cloudera Manager

Azul Open JDK 8 and 11 are not supported with Cloudera Manager. To use Azul Open JDK 8 or 11 for Cloudera Manager RPM/DEBs, you must manually create a symlink between the Zulu JDK installation path and the default JDK path.

After installing Azul Open JDK8 or 11, you must run the following commands on all the hosts in the cluster:

Azul Open JDK 8

RHEL or SLES

```
# sudo ln -s /usr/lib/jvm/java-8-zulu-openjdk-jdk /usr/lib/jvm/j
ava-8-openjdk
```

Ubuntu or Debian

```
# sudo ln -s /usr/lib/jvm/zulu-8-amd64 /usr/lib/jvm/java-8-openj
dk
```

Azul Open JDK 11

For DEBs only

sudo ln -s /usr/lib/jvm/zulu-11-amd64 /usr/lib/jvm/java-11

CDPD-62834: Status of the deleted table is seen as ACTIVE in Atlas after the completion of navigator2atlas migration process

The status of the deleted table displays as ACTIVE.

None

CDPD-62837: During the navigator2atlas process, the hive_storagedesc is incomplete in Atlas

For the hive_storagedesc entity, some of the attributes are not getting populated.

None

OPSAPS-69897: NPE in Ozone replication from CM 7.7.1 to CM 7.11.3

When you use source Cloudera Manager 7.7.1 and target Cloudera Manager 7.11.3 for Ozone replication policies, the policies fail with Failure during PreOzoneCopyListingCheck execution: null error. This is because the target Cloudera Manager 7.11.3 does not retrieve the required source bucket information for validation from the source Cloudera Manager 7.7.1 during the PreCopyListingCheck command phase. You come across this error when you use source Cloudera Manager versions lower than 7.10.1 and target Cloudera Manager versions higher than or equal to 7.10.1 in an Ozone replication policy.

Upgrade the source Cloudera Manager to 7.11.3 or higher version.

OPSAPS-69481: Some Kafka Connect metrics missing from Cloudera Manager due to conflicting definitions

The metric definitions for kafka_connect_connector_task_metrics_batch_size_avg and kafka_connect_connector_task_metrics_batch_size_max in recent Kafka CSDs conflict with previous definitions in other CSDs. This prevents Cloudera Manager from registering these metrics. It also results in SMM returning an error. The metrics also cannot be monitored in Cloudera Manager chart builder or queried using the Cloudera Manager API.

Contact Cloudera support for a workaround.

Following are the list of fixed issues that were shipped for Cloudera Manager 7.11.3 CHF3 (version: 7.11.3.4-50275000):

OPSAPS-69267: Extend Java opts for Impala to support JDK17 + Isilon

Impala no longer reports IllegalAccessErrors on sun.net.dns with Java 17 and Isilon.

OPSAPS-69485: Invalid mapred-site.xml due to double dash in comments

The string '--' is not allowed in XML comments. Cloudera Manager incorporates values from the safety valve into XML comments. Therefore, XML configuration file generation fails if the safety valve contains '--'.

Cloudera Manager replaces the '--' characters in XML configuration file comments with — which is the Unicode character of '--'.

CDPD-62464: Java process called by navatlas.sh tool fails on JDK-8 version

While running nav2atlas.sh script on OracleJDK 8 an error message is thrown and returns code 0 on an unsuccessful run.

OPSAPS-69340: Dlog4j.configurationFile annotation is not working with the log4j library of the Cloudera Manager Server.

The incorrect notation used in defining the log4j configuration file name (which is Dlog4j.configura tionFile annotation) is preventing the Cloudera Manager Server from receiving updates made to thelog4j.properties file. This issue is fixed now.

OPSAPS-69022: Rack Topology is not updated on Ozone DataNode.

Ozone, Kudu, and Cruise Control are rack aware services but topology mapping for the hosts containing roles from these services only are not updated in previous versions unless an HDFS or YARN role was present on these hosts. Fixed this issue in Cloudera Manager. Hosts containing Ozone, Kudu, and Cruise Control roles should now get the right topology mapping regardless of other roles present on the host.

OPSAPS-69414: Expose missing Ozone metrics to Cloudera Manager.

Three new Ozone metrics (number of datanodes, total capacity, and used space) are exposed to Cloudera Manager.

OPSAPS-69063: Concurrent policy creation to multiple targets

Sometimes, standard error or standard output retrieval of Cloudera Manager commands would fail because of a Java-related issue which affected the HTTPS connections using TLSv1.3 protocol. This resulted in different failures when the HBase replication commands were run remotely from the destination cluster on the source cluster. This issue is now resolved.

OPSAPS-69257: Zeppelin: Interpreter logs are not getting printed

A new parameter is added in SDL to pass the zeppelin log file name dynamically to the log4.properties file. This allows Zeppelin to log the interpreter logs in the CM Cluster.

OPSAPS-69131: Unable to install Zeppelin on CDP 7.1.9 with Spark 3 on RHEL 9

Zeppelin installation was failing on RHEL 9 because Spark 2 was configured as a mandatory dependency in the Zeppelin SDL file. Spark 2 could not be installed on RHEL 9 due to the python version incompatibility. This issue is now resolved.

OPSAPS-69194: JDK 17 support for HBase Indexer

This fix makes the KS-Indexer service JDK 17 compatible.

OPSAPS-69378: Increase default certificate lifetime

Default behaviour intact. Default expiry time - 1 year for all certs issued by RKE CA (except DB since it is not issued by RKE CA) Expiry time can be adjusted via CM parameter cluster_signing_duration. To apply the changes - Rolling restart for ECS and Rotate of Vault, DB, ecs webhook, ingress certs is required.

OPSAPS-69329: Configure higher 'worker-shutdown-timeout' value for nginx ingress controller

Updated rke-nginx-ingress-controller worker-shutdown-timeout config to 24 hrs.

OPSAPS-69250: ECS Server restart failed as it requires "yum install" from repo

The "yum install" line has been removed from the ECS Server startup script.

CDPD-62464: Java process called by navatlas.sh tool fails on JDK-8 version

Explicitly added eclipse-collections dependencies of version which is compatible with JDK-8 version.

OPSAPS-69245: Oozie issue in FIPS environments

In FIPS environment, Oozie needed the FIPS-related Java options to be present in HADOOP_CLIENT_OPTS in order to pick them up. Java options also needed to be added to container localizers. An admin option parameter for the localizers to pass these options has been created and will not to interfere with the user-defined options. This issue is now fixed.

OPSAPS-69406: Existing HDFS and HBase snapshot policy configuration can be edited

The **Edit Configuration** modal window appears when you click Actions Edit Configuration on the Cloudera Manager Replication Snapshot Policies page for existing HDFS or HBase snapshot policies.

The repositories for Cloudera Manager 7.11.3-CHF3 are listed in the following table:

Table 3: Cloudera Manager 7.11.3-CHF3

Repository Type	Repository Location
RHEL 9 Compatible	Repository:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.4-50275000/redhat9/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.4-50275000/redhat9/yum/cloudera-manager.repo
RHEL 8 Compatible	Repository:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.4-50275000/redhat8/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.4-50275000/redhat8/yum/cloudera-manager.repo

Repository Type	Repository Location
RHEL 7 Compatible	Repository:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.4-50275000/redhat7/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.4-50275000/redhat7/yum/cloudera-manager.repo
SLES 15	Repository:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.4-50275000/sles15/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/cm7/patc h/7.11.3.4-50275000/sles15/yum/cloudera-manager.repo
SLES 12	Repository:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.4-50275000/sles12/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/cm7/patc h/7.11.3.4-50275000/sles12/yum/cloudera-manager.repo
Ubuntu 20	Repository:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.4-50275000/ubuntu2004/apt
	Repository File:
	https://username:password@archive.cloudera.com/p/cm7/patc h/7.11.3.4-50275000/ubuntu2004/apt/cloudera-manager.list

Cloudera Manager 7.11.3 Cumulative hotfix 2

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 2.

This cumulative hotfix was released on December 21, 2023.



Note: Contact Cloudera Support for questions related to any specific hotfixes.

New features and changed behavior for Cloudera Manager 7.11.3 CHF2 (version: 7.11.3.3-47960007): Replicate Hive ACID tables and Iceberg tables in Dell EMC Isilon storage clusters using Hive ACID table replication policies and Iceberg replication policies respectively

You can replicate Hive ACID tables and Iceberg tables, using replication policies, between CDP Private Cloud Base 7.1.9 or higher clusters on Dell EMC Isilon storage using Cloudera Manager 7.11.3 CHF2 or higher versions.

Wait timeout for regenerating credentials in Active Directory (AD)

Cloudera Manager supports a new parameter ad_wait_time_for_regenerate to indicate the wait timeout period after deleting an old principal and allowing this deletion to replicate on all AD servers in sufficient period of time. This ensures a successful creation of a new principal after the deletion process. Set the timeout value according to your AD setup (number of AD server replicas). If the timeout value is too low, then an error of *stale principal* might occur (ldap_add: Constraint violation (19) additional info: 000021C8: AtrErr: DSID-03200EB7, #1: 0: 000021C8: DSID-03200EB7, problem 1005 (CONSTRAINT_ATT_TYPE), data 0, Att 90290 (userPrincipalName)).



Important:

This parameter will not be operational if the value is set to 0.

You might set this parameter before running the following commands:

- Generate Missing Credentials
- Regenerate Selected

FIPS support for JDK11 in Kudu

Added FIPS support for JDK11 in Kudu.

FIPS support for JDK11 in Hive

Added the required JVM arguments in Hive processes in order to execute on a FIPS enabled cluster.

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF2 (version: 7.11.3.3-47960007):

OPSAPS-69406: Cannot edit existing HDFS and HBase snapshot policy configuration

The **Edit Configuration** modal window does not appear when you click Actions Edit Configuration on the Cloudera Manager Replication Snapshot Policies page for existing HDFS or HBase snapshot policies.

None.

OPSAPS-68340: Zeppelin paragraph execution fails with the User not allowed to impersonate error.

Starting from Cloudera Manager 7.11.3, Cloudera Manager auto-configures the livy_admin_users configuration when Livy is run for the first time. If you add Zeppelin or Knox services later to the existing cluster and do not manually update the service user, the User not allowed to impersonate error is displayed.

If you add Zeppelin or Knox services later to the existing cluster, you must manually add the respective service user to the livy_admin_users configuration in the Livy configuration page.

OPSAPS-69342: Access issues identified in MariaDB 10.6 were causing discrepancies in High Availability (HA) mode

MariaDB 10.6, by default, includes the property require_secure_transport=ON in the configuration file (/etc/my.cnf), which is absent in MariaDB 10.4. This setting prohibits non-TLS connections, leading to access issues. This problem is observed in High Availability (HA) mode, where certain operations may not be using the same connection.

To resolve the issue temporarily, you can either comment out or disable the line require_secure_t ransport in the configuration file located at /etc/my.cnf.

OPSAPS-68452: Azul Open JDK 8 and 11 are not supported with Cloudera Manager

Azul Open JDK 8 and 11 are not supported with Cloudera Manager. To use Azul Open JDK 8 or 11 for Cloudera Manager RPM/DEBs, you must manually create a symlink between the Zulu JDK installation path and the default JDK path.

After installing Azul Open JDK8 or 11, you must run the following commands on all the hosts in the cluster:

Azul Open JDK 8

RHEL or SLES

```
# sudo ln -s /usr/lib/jvm/java-8-zulu-openjdk-jdk /usr/lib/jvm/j
ava-8-openjdk
```

Ubuntu or Debian

```
# sudo ln -s /usr/lib/jvm/zulu-8-amd64 /usr/lib/jvm/java-8-openj
dk
```

Azul Open JDK 11

For DEBs only

```
# sudo ln -s /usr/lib/jvm/zulu-11-amd64 /usr/lib/jvm/java-11
```

OPSAPS-69340: Dlog4j.configurationFile annotation is not working with the log4j library of the Cloudera Manager Server.

The incorrect notation used in defining the log4j configuration file name (which is Dlog4j.configura tionFile annotation) is preventing the Cloudera Manager Server from receiving updates made to thelog4j.properties file.

Perform the following steps:

1. Edit the /etc/default/cloudera-scm-server file by adding the following line:

```
export CMF_JAVA_OPTS="-Dlog4j.configuration=file:/etc/cloude
ra-scm-server/log4j.properties $CMF_JAVA_OPTS"
```

2. Restart the Cloudera Manager Server by running the following command:

sudo systemctl restart cloudera-scm-server

CDPD-62464: Java process called by navatlas.sh tool fails on JDK-8 version

While running nav2atlas.sh script on OracleJDK 8 an error message is thrown and returns code 0 on an unsuccessful run.

You must install JDK-11 version on the host. Make sure not to put into the default path and JAVA _HOME. In a shell, set the JAVA_HOME to this location and run the nav2atlas.sh script.

CDPD-62834: Status of the deleted table is seen as ACTIVE in Atlas after the completion of navigator2atlas migration process

The status of the deleted table displays as ACTIVE.

None

CDPD-62837: During the navigator2atlas process, the hive_storagedesc is incomplete in Atlas

For the hive_storagedesc entity, some of the attributes are not getting populated.

None

OPSAPS-69897: NPE in Ozone replication from CM 7.7.1 to CM 7.11.3

When you use source Cloudera Manager 7.7.1 and target Cloudera Manager 7.11.3 for Ozone replication policies, the policies fail with Failure during PreOzoneCopyListingCheck execution: null error. This is because the target Cloudera Manager 7.11.3 does not retrieve the required source bucket information for validation from the source Cloudera Manager 7.7.1 during the PreCopyListingCheck command phase. You come across this error when you use source Cloudera Manager versions lower than 7.10.1 and target Cloudera Manager versions higher than or equal to 7.10.1 in an Ozone replication policy.

Upgrade the source Cloudera Manager to 7.11.3 or higher version.

OPSAPS-69481: Some Kafka Connect metrics missing from Cloudera Manager due to conflicting definitions

The metric definitions for kafka_connect_connector_task_metrics_batch_size_avg and kafka_connect_connector_task_metrics_batch_size_max in recent Kafka CSDs conflict with previous definitions in other CSDs. This prevents Cloudera Manager from registering these metrics. It also results in SMM returning an error. The metrics also cannot be monitored in Cloudera Manager chart builder or queried using the Cloudera Manager API.

Contact Cloudera support for a workaround.

Following are the list of fixed issues that were shipped for Cloudera Manager 7.11.3 CHF2 (version: 7.11.3.3-47960007):

OPSAPS-68689: Unable to emit the LDAP Bind password in core-site.xml for client configurations

If the CDP cluster has LDAP group to OS group mapping enabled, then applications running in Spark or Yarn would fail to authenticate to the LDAP server when trying to use the LDAP bind account during the LDAP group search.

This is because the LDAP bind password was not passed to the /etc/hadoop/conf/core-site.xml file. This was intended behavior to prevent leaking the LDAP bind password in a clear text field.

To fix this issue, perform the instructions from the Emitting the LDAP Bind password in coresite.xml for client configurations section to emit the LDAP Bind password in core-site.xml for client configurations.

OPSAPS-60139: Staleness performance issue in clusters with a large number of roles

In large clusters, Cloudera Manager takes a long time to display the Configuration Staleness icon after a service configuration change. This issue is fixed now by improving the performance of the staleness-checking algorithm.

OPSAPS-68722: Java heap size can now be configured

You can now customize Java heap size in YARN Queue Manager. Although the default for this setting should be valid in most deployment scenarios, you have the option to update the setting only if a given cluster has run into memory-management issues, otherwise, the settings can remain.

OPSAPS-68217: Add post replication diff to compare files

You can now trace files that go missing during snapshot-based cloud replication. To trace and debug the issue, perform the following steps:

1. Go to the Clusters HDFS Configuration tab.

- 2. To enable the debug steps, complete the following steps:
 - **a.** Search for the HDFS Replication Environment Advanced Configuration Snippet (Safety Valve) property.
 - **b.** Add the following key-value pair, and Save the changes:

SCHEDULES_WITH_ADDITIONAL_DEBUG_STEPS = [***comma-separated list of numerical IDs of all the applicable replication policies***]

- c. Search for the HDFS Replication Advanced Configuration Snippet (Safety Valve) for hdfssite.xml property.
- d. Add the following key-value pair, and Save the changes:

com.cloudera.enterprise.distcp.post-copy-reconciliation.fail-on = MISSING_ON_TARGET

The possible values for this parameter include MISSING_ON_SOURCE, MISSING_ON_TARGET, MISSING_ON_BOTH, ANY_MISSING, and NONE. The default is NONE.



Note: The key-value pair fails the replication job if there are any missing files on the target. This failure does not invalidate the snapshot, which means that if a file is missing, it remains missing until you manually force-run a bootstrap replication which is an optional action. If you do not set the key-value pair, the replication job continues and generates the debug info.

- 3. To enable extra logging, complete the following steps:
 - **a.** Search for the HDFS Replication Environment Advanced Configuration Snippet (Safety Valve) property.
 - b. Add the following key-value pair, and Save the changes:

EXTRA_LOG_CONFIGS_\$SCHEDULE_ID =

```
log4j.rootLogger=INFO,console;
hadoop.root.logger=INFO,console;log4j.appender.console=or
g.apache.log4j.ConsoleAppender;
log4j.appender.console.target=System.err;log4j.appender.cons
ole.layout=org.apache.log4j.PatternLayout;
log4j.appender.console.layout.ConversionPattern=%d{yy/MM/dd
HH:mm:ss} %p %c{2}: %m%n;
log4j.logger.org.apache.hadoop.fs.azurebfs.services.AbfsIoU
tils=DEBUG,console;
log4j.logger.org.apache.hadoop.fs.azurebfs.services.AbfsClie
nt=DEBUG,console;
log4j.logger.distcp.SimpleCopyListing=DEBUG,console;log4j.
logger.distcp.SnapshotDiffGenerator=DEBUG,console
```



Note: Replace \$SCHEDULE_ID with the numerical ID of the replication policy this should apply to.

The extra debug logs are collected on HDFS in the \$logDir/debug directory. For example, the log location hdfs://user/hdfs/.cm/distcp/2023-08-24_206/debug

OPSAPS-68855: Fix replication policy deletion for Hive ACID replication policies using Dell Powerscale Isilon clusters

The Hive ACID replication policy can be deleted successfully on CDP Private Cloud Base 7.1.9 or higher clusters with Dell EMC Isilon storage using Cloudera Manager 7.11.3 CHF2 or higher versions.

OPSAPS-68516: Ozone replication diagnostic bundle collection

Replication Manager generates diagnostic information bundle for Ozone replication policies.

OPSAPS-68698: Replication command type is incorrectly reported for Ozone incremental replication

When you create an Ozone replication policy using the "Incremental with fallback to full file listing" Listing type, the Ozone replication command type correctly reports the file listing type for the run.

The first run of an Ozone replication policy creates a snapshot during the run, but it cannot calculate a snapshot diff because there is no previous snapshot. In this case, full file listing is used for the first run of the policy. This is now reported correctly as FULL_FILE_LISTING_FALLBACK.

OPSAPS-68856: Fix Hive ACID replication policy creation when using Dell Powerscale Isilon clusters

The Hive ACID replication policy can be created successfully on CDP Private Cloud Base 7.1.9 or higher clusters with Dell EMC Isilon storage using Cloudera Manager 7.11.3 CHF2 or higher versions.

OPSAPS-68995: Convert some DistCp feature checks from CM version checks to feature flags

To ensure interoperability between different cumulative hotfixes (CHF), the NUM_FETCH_THREADS, DELETE_LATEST_SOURCE_SNAPSHOT_ON_JOB_FAILURE, and RAISE_SNAPSHOT_DIFF_FAILURES DistCp features must be published as feature flags.

OPSAPS-68658: Source ozone service ID is used as target

Ozone replication policies do not fail when the Ozone service name is different on the source and destination clusters because Ozone replication uses the destination Ozone service name during the path normalization process.

OPSAPS-68526 - Iceberg Replication support for Dell Powerscale

Iceberg replication policies run successfully on Kerberos-enabled clusters on Dell EMC Isilon storage. For more information, see Adding custom Kerberos keytab and Kerberos principal for replication policies.

The repositories for Cloudera Manager 7.11.3-CHF2 are listed in the following table:

Table 4: Cloudera Manager 7.11.3-CHF2

Repository Type	Repository Location
RHEL 9 Compatible	Repository:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.3-47960007/redhat9/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.3-47960007/redhat9/yum/cloudera-manager.repo
RHEL 8 Compatible	Repository:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.3-47960007/redhat8/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.3-47960007/redhat8/yum/cloudera-manager.repo

Repository Type	Repository Location
RHEL 7 Compatible	Repository:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.3-47960007/redhat7/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.3-47960007/redhat7/yum/cloudera-manager.repo
SLES 15	Repository:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.3-47960007/sles15/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/cm7/patc h/7.11.3.3-47960007/sles15/yum/cloudera-manager.repo
SLES 12	Repository:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.3-47960007/sles12/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/cm7/patc h/7.11.3.3-47960007/sles12/yum/cloudera-manager.repo
Ubuntu 20	Repository:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.3-47960007/ubuntu2004/apt
	Repository File:
	https://username:password@archive.cloudera.com/p/cm7/patc h/7.11.3.3-47960007/ubuntu2004/apt/cloudera-manager.list
IBM PowerPC RHEL 9	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.3-47960007/redhat9-ppc/yum
IBM PowerPC RHEL 8	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.3-47960007/redhat8-ppc/yum

Cloudera Manager 7.11.3 Cumulative hotfix 1

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 1.

This cumulative hotfix was released on November 2, 2023.



Note: Contact Cloudera Support for questions related to any specific hotfixes.

New features and changed behavior for Cloudera Manager 7.11.3 CHF1 (version: 7.11.3.2-46642574): Cloudera Navigator role instances under the Cloudera Management Service are no longer available while using Cloudera Runtime 7.1.9 CHF1 version.

You must first migrate Cloudera Navigator to Atlas before you upgrade from CDH 6.x + ClouderaManager 6.x / 7.x to CDP 7.1.9 CHF1 + Cloudera Manager 7.11.3 Latest cumulative hotfix. For more information, you must refer to Migrating from Cloudera Navigator to Atlas using Cloudera Manager 6 and Migrating from Cloudera Manager to Atlas using Cloudera Manager 7.

OpenJDK 17 (TCK certified) support for the Cloudera Manager 7.11.3 CHF1 and operating systems

Cloudera Manager 7.11.3 CHF1 now supports OpenJDK 17 (TCK certified) on RHEL 7, RHEL 8, RHEL 9, Ubuntu 20, and SLES 15.

You must upgrade to Cloudera Manager 7.11.3 CHF1 or higher, before upgrading to OpenJDK 17 (TCK certified).

Replicate Hive external tables in Dell EMC Isilon storage clusters using Hive external table replication policies

You can use Hive external table replication policies in CDP Private Cloud Base Replication Manager to replicate Hive external tables between Dell EMC Isilon storage clusters where the 7.1.9 clusters use Cloudera Manager 7.11.3 CHF1 or higher versions.

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF1 (version: 7.11.3.2-46642574):

OPSAPS-69406: Cannot edit existing HDFS and HBase snapshot policy configuration

The **Edit Configuration** modal window does not appear when you click Actions Edit Configuration on the Cloudera Manager Replication Snapshot Policies page for existing HDFS or HBase snapshot policies.

None.

OPSAPS-68340: Zeppelin paragraph execution fails with the User not allowed to impersonate error.

Starting from Cloudera Manager 7.11.3, Cloudera Manager auto-configures the livy_admin_users configuration when Livy is run for the first time. If you add Zeppelin or Knox services later to the existing cluster and do not manually update the service user, the User not allowed to impersonate error is displayed.

If you add Zeppelin or Knox services later to the existing cluster, you must manually add the respective service user to the livy_admin_users configuration in the Livy configuration page.

OPSAPS-68689: Unable to emit the LDAP Bind password in core-site.xml for client configurations

If the CDP cluster has LDAP group to OS group mapping enabled, then applications running in Spark or Yarn would fail to authenticate to the LDAP server when trying to use the LDAP bind account during the LDAP group search.

This is because the LDAP bind password was not passed to the /etc/hadoop/conf/core-site.xml file. This was intended behavior to prevent leaking the LDAP bind password in a clear text field.

Set the LDAP Bind password through the HDFS client configuration safety valve.

- **1.** On the Cloudera Manager UI, navigate to the HDFS service, by clicking on the HDFS service under the Cluster.
- **2.** Click the Configuration tab. Search for the HDFS Client Advanced Configuration Snippet (Safety Valve) for hdfs-site.xml configuration parameter.

- **3.** Add an entry with the following values:
 - Name = hadoop.security.group.mapping.ldap.bind.password
 - Value = (Enter the LDAP bind password here)
 - Description = Password for LDAP bind account
- 4. Then click the Save Changes button to save the safety valve entry.
- **5.** Perform the instructions from the Manually Redeploying Client Configuration Files to manually deploy client configuration files to the cluster.

OPSAPS-69342: Access issues identified in MariaDB 10.6 were causing discrepancies in High Availability (HA) mode

MariaDB 10.6, by default, includes the property require_secure_transport=ON in the configuration file (/etc/my.cnf), which is absent in MariaDB 10.4. This setting prohibits non-TLS connections, leading to access issues. This problem is observed in High Availability (HA) mode, where certain operations may not be using the same connection.

To resolve the issue temporarily, you can either comment out or disable the line require_secure_t ransport in the configuration file located at /etc/my.cnf.

OPSAPS-68452: Azul Open JDK 8 and 11 are not supported with Cloudera Manager

Azul Open JDK 8 and 11 are not supported with Cloudera Manager. To use Azul Open JDK 8 or 11 for Cloudera Manager RPM/DEBs, you must manually create a symlink between the Zulu JDK installation path and the default JDK path.

After installing Azul Open JDK8 or 11, you must run the following commands on all the hosts in the cluster:

Azul Open JDK 8

RHEL or SLES

```
# sudo ln -s /usr/lib/jvm/java-8-zulu-openjdk-jdk /usr/lib/jvm/j
ava-8-openjdk
```

Ubuntu or Debian

```
# sudo ln -s /usr/lib/jvm/zulu-8-amd64 /usr/lib/jvm/java-8-openj
dk
```

Azul Open JDK 11

For DEBs only

```
# sudo ln -s /usr/lib/jvm/zulu-11-amd64 /usr/lib/jvm/java-11
```

OPSAPS-69481: Some Kafka Connect metrics missing from Cloudera Manager due to conflicting definitions

The metric definitions for kafka_connect_connector_task_metrics_batch_size_avg and kafka_connect_connector_task_metrics_batch_size_max in recent Kafka CSDs conflict with previous definitions in other CSDs. This prevents Cloudera Manager from registering these metrics. It also results in SMM returning an error. The metrics also cannot be monitored in Cloudera Manager chart builder or queried using the Cloudera Manager API.

Contact Cloudera support for a workaround.

OPSAPS-68559: On-premises to on-premises Hive external replication won't work with a cloud target

You cannot replicate Hive external tables using Hive external table replication policies from an onpremises cluster to another on-premises cluster with an external account to replicate the Hive data only to the cloud.

None

OPSAPS-68658: Source ozone service id used as target

Ozone replication policies fail when the Ozone service ID is different on the source and destination clusters because Ozone replication uses the destination Ozone service ID during the path normalization process.

None

OPSAPS-68698: Replication command type is incorrectly reported for Ozone incremental replications

When you create an Ozone replication policy using "Incremental only" or "Incremental with fallback to full file listing" Listing types, sometimes the Ozone replication command type is incorrectly reported for different types of runs.

None

OPSAPS-42908: "User:hdfs not allowed to do DECRYPT_EEK" error appears for Hive external table replication policies

When you run Hive external table replication policies on clusters using Ranger KMS, the "User:hdfs not allowed to do 'DECRYPT_EEK'" error appears when you do not use the hdfs username.

Edit the Hive external table replication policy, and configure the Advanced Directory for metadata file field to a new directory that is not encrypted. The replication policy uses this directory to store the transient data during Hive replication.

OPSAPS-69897: NPE in Ozone replication from CM 7.7.1 to CM 7.11.3

When you use source Cloudera Manager 7.7.1 and target Cloudera Manager 7.11.3 for Ozone replication policies, the policies fail with Failure during PreOzoneCopyListingCheck execution: null error. This is because the target Cloudera Manager 7.11.3 does not retrieve the required source bucket information for validation from the source Cloudera Manager 7.7.1 during the PreCopyListingCheck command phase. You come across this error when you use source Cloudera Manager versions lower than 7.10.1 and target Cloudera Manager versions higher than or equal to 7.10.1 in an Ozone replication policy.

Upgrade the source Cloudera Manager to 7.11.3 or higher version.

CDPD-62464: Java process called by navatlas.sh tool fails on JDK-8 version

While running nav2atlas.sh script on OracleJDK 8 an error message is thrown and returns code 0 on an unsuccessful run.

You must install JDK-11 version on the host. Make sure not to put into the default path and JAVA _HOME. In a shell, set the JAVA_HOME to this location and run the nav2atlas.sh script.

CDPD-62834: Status of the deleted table is seen as ACTIVE in Atlas after the completion of navigator2atlas migration process

The status of the deleted table displays as ACTIVE.

None

CDPD-62837: During the navigator2atlas process, the hive_storagedesc is incomplete in Atlas

For the hive_storagedesc entity, some of the attributes are not getting populated.

None

Following are the list of fixed issues that were shipped for Cloudera Manager 7.11.3 CHF1 (version: 7.11.3.2-46642574):

OPSAPS-68664: Added the support for JDK 17 in HDFS.

This issue is resolved.

OPSAPS-68550: Ozone Canary failing with unknown option --skipTrash.

This issue is resolved.

OPSAPS-66023: Error message about an unsupported ciphersuite while upgrading or installing cluster with the latest FIPS compliance

When upgrading or installing a FIPS enabled cluster, Cloudera Manager is unable to download the new CDP parcel from the Cloudera parcel archive.

Cloudera Manager displays the following error message:

HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA

This issue is fixed now by correcting the incorrect ciphersuite selection.

OPSAPS-65504: Upgraded Apache Ivy version

The Apache Ivy version is upgraded from 2.x.x to 2.5.1 version to fix CVE issues.

OPSAPS-68500: The cloudera-manager-installer.bin fails to reach Ubuntu 20 repository on the Archive URL due to redirections

Agent Installation with Cloudera Manager on Ubuntu20 platform does not function when the selfinstaller method (using the installer.bin file) is employed to install Cloudera Manager. The failure mode is that Cloudera Manager Agent installation step will fail with an error message saying "The repository 'https://archive.cloudera.com/p/cm7/7.11.3/ubuntu2004/apt focal-cm7 InRelease' is not signed."

This issue is fixed now.

OPSAPS-68422: Incorrect HBase shutdown command can lead to inconsistencies

Cloudera Manager uses an incomplete stop command when you stop the HBase service or the corresponding roles on a 7.1.8 or higher private cloud cluster. Due to this, the Cloudera Manager cannot gracefully stop the processes and kill them after a set timeout. This could lead to metadata corruption.

This issue is fixed now.

OPSAPS-68506: Knox CSD changes for readiness check

A readiness endpoint was added to determine whether Knox is ready to receive traffic. Clouder Manager checks the state of Knox after startup to reduce downtime during rolling restarts.

OPSAPS-68424 Impala: CM Agent unable to extract logs to TP export directory

Impala queries were not displaying in the Cloudera Observability and Workload XM web User Interfaces. This was due to an internal error that was stopping Cloudera Manager from pushing the Impala profile to the Telemetry Publisher logs directory.

This Issue is now fixed and the Telemetry Publisher's log extraction has been re-enable.



Note: If you are using a Cloudera Manager version between 7.11.2.0 and 7.11.3, Cloudera recommends upgrading to Cloudera Manager 7.11.3.CHF1.

OPSAPS-68697 Error while generating email template resulting in an inability to trigger mail notification

Cloudera Observability and Workload XM were unable to trigger an email notification when an Impala query matched the Auto Action's alert threshold value.

This problem occurred when both the following conditions were met:

- The Auto Action is triggered for an Impala Scope.
- The length of the Impala query on which the action event is triggered is less than 36 characters.

This issue is now fixed.

OPSAPS-68798 Auto Actions not using proxy while connecting to DBUS

The Cloudera Observability and Workload XM Auto Actions feature was not recognizing the proxy server credentials, even when they were correct and the proxy server was enabled in Telemetry Publisher.

This issue is now fixed.



Note: Users must ensure they have the correct proxy server setup and that the proxy server is enabled in Telemetry Publisher.

Known issue: OPSAPS-68629: HDFS HTTPFS GateWay is not able to start with custom krb5.conf location set in Cloudera Manager.

On a cluster with a custom krb5.conf file location configured in Cloudera Manager, HDFS HTTPFS role is not able to start because it does not have the custom Kerberos configuration file setting properly propagated to the service, and therefore it fails with a Kerberos related exception: in threa d "main" java.io.IOException: Unable to initialize WebAppContext at org.apache.hadoop.http .HttpServer2.start(HttpServer2.java:1240) at org.apache.hadoop.fs.http.server.HttpFSServerWebS erver.start(HttpFSServerWebServer.java:131) at org.apache.hadoop.fs.http.server.HttpFSServer WebServer.main(HttpFSServerWebServer.java:162) Caused by: java.lang.IllegalArgumentExc eption: Can't get Kerberos realm at org.apache.hadoop.security.HadoopKerberosName.setConfi guration(HadoopKerberosName.java:71) at org.apache.hadoop.security.UserGroupInformation. initialize(UserGroupInformation.java:329) at org.apache.hadoop.security.UserGroupInforma tion.setConfiguration(UserGroupInformation.java:380) at org.apache.hadoop.lib.service.hadoop .FileSystemAccessService.init(FileSystemAccessService.java:166) at org.apache.hadoop.lib.s erver.BaseService.init(BaseService.java:71) at org.apache.hadoop.lib.server.Server.initService s(Server.java:581) at org.apache.hadoop.lib.server.Server.init(Server.java:377) at org.apache.ha doop.fs.http.server.HttpFSServerWebApp.init(HttpFSServerWebApp.java:100) at org.apache.h adoop.lib.servlet.ServerWebApp.contextInitialized(ServerWebApp.java:158) at org.eclipse.jett y.server.handler.ContextHandler.callContextInitialized(ContextHandler.java:1073) at org.ecli pse.jetty.servlet.ServletContextHandler.callContextInitialized(ServletContextHandler.java:572) org.eclipse.jetty.server.handler.ContextHandler.contextInitialized(ContextHandler.java:1002) at

at org.eclipse.jetty.servlet.ServletHandler.initialize(ServletHandler.java:765) at org.eclipse.jetty. servlet.ServletContextHandler.startContext(ServletContextHandler.java:379) at org.eclipse.jett y.webapp.WebAppContext.startWebapp(WebAppContext.java:1449) at org.eclipse.jetty.webap p.WebAppContext.startContext(WebAppContext.java:1414) at org.eclipse.jetty.server.handler. ContextHandler.doStart(ContextHandler.java:916) at org.eclipse.jetty.servlet.ServletContextHan dler.doStart(ServletContextHandler.java:288) at org.eclipse.jetty.webapp.WebAppContext.doStart (WebAppContext.java:524) at org.eclipse.jetty.util.component.AbstractLifeCycle.start(AbstractL ifeCycle.java:73) at org.eclipse.jetty.util.component.ContainerLifeCycle.start(ContainerLifeCycle. java:169) at org.eclipse.jetty.util.component.ContainerLifeCycle.doStart(ContainerLifeCycle. java:117) at org.eclipse.jetty.server.handler.AbstractHandler.doStart(AbstractHandler.java:97) a t org.eclipse.jetty.util.component.AbstractLifeCycle.start(AbstractLifeCycle.java:73) at org.ecl ipse.jetty.util.component.ContainerLifeCycle.start(ContainerLifeCycle.java:169) at org.eclipse.jet ty.server.Server.start(Server.java:423) at org.eclipse.jetty.util.component.ContainerLifeCycle.doS tart(ContainerLifeCycle.java:110) at org.eclipse.jetty.server.handler.AbstractHandler.doStart(Abst ractHandler.java:97) at org.eclipse.jetty.server.Server.doStart(Server.java:387) at org.eclipse.jett y.util.component.AbstractLifeCycle.start(AbstractLifeCycle.java:73) at org.apache.hadoop.http. HttpServer2.start(HttpServer2.java:1218) ... 2 more Caused by: java.lang.IllegalArgumentExcept ion: KrbException: Cannot locate default realm at java.security.jgss/javax.security.auth.kerberos. KerberosPrincipal.<init>(KerberosPrincipal.java:174) at org.apache.hadoop.security.authentic ation.util.KerberosUtil.getDefaultRealm(KerberosUtil.java:108) at org.apache.hadoop.security .HadoopKerberosName.setConfiguration(HadoopKerberosName.java:69)

- 1. Log in to Cloudera Manager.
- 2. Select the HDFS service.
- **3.** Select Configurations tab.
- 4. Search for HttpFS Environment Advanced Configuration Snippet (Safety Valve)

- **5.** Add to or extend the HADOOP_OPTS environment variable with the following value: Djava.security.krb5.conf=<the custom krb5.conf location>
- 6. Click Save Changes.

The repositories for Cloudera Manager 7.11.3-CHF1 are listed in the following table:

Table 5: Cloudera Manager 7.11.3-CHF1

Repository Type	Repository Location
RHEL 9 Compatible	Repository:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.2-46642574/redhat9/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.2-46642574/redhat9/yum/cloudera-manager.repo
RHEL 8 Compatible	Repository:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.2-46642574/redhat8/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.2-46642574/redhat8/yum/cloudera-manager.repo
RHEL 7 Compatible	Repository:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.2-46642574/redhat7/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.2-46642574/redhat7/yum/cloudera-manager.repo
SLES 15	Repository:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.2-46642574/sles15/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/cm7/patc h/7.11.3.2-46642574/sles15/yum/cloudera-manager.repo
SLES 12	Repository:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.2-46642574/sles12/yum
	Repository File:
	https://username:password@archive.cloudera.com/p/cm7/patc h/7.11.3.2-46642574/sles12/yum/cloudera-manager.repo

Repository Type	Repository Location
Ubuntu 20	Repository:
	https://username:password@archive.cloudera.com/p/cm7/ patch/7.11.3.2-46642574/ubuntu2004/apt
	Repository File:
	https://username:password@archive.cloudera.com/p/cm7/patc h/7.11.3.2-46642574/ubuntu2004/apt/cloudera-manager.repo
IBM PowerPC RHEL 8	https://username:password@archive.cloudera.com/p/cm7/patc h/7.11.3.2-46642574/redhat8-ppc/yum
IBM PowerPC RHEL 9	https://username:password@archive.cloudera.com/p/cm7/patc h/7.11.3.2-46642574/redhat9-ppc/yum