

Cloudera Manager 7.13.1

Unified Cloudera Manager Release Notes

Date published: 2024-05-15

Date modified: 2024-05-15

CLOUDERA

<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

- Cloudera Manager 7.13.1 Release Notes..... 4**
 - What's New in Cloudera Manager 7.13.1.....4
 - What's new in Platform Support..... 5
 - Fixed Issues in Cloudera Manager 7.13.1..... 6
 - Known Issues in Cloudera Manager 7.13.1..... 12
 - Behavioral Changes In Cloudera Manager 7.13.1..... 17
 - Fixed Common Vulnerabilities and Exposures in Cloudera Manager 7.13.1.....17
 - Deprecation notices in Cloudera Manager 7.13.1.....21
 - Platform and OS..... 21

Cloudera Manager 7.13.1 Release Notes

You can review the Release Notes of Cloudera Manager 7.13.1 associated with unified Cloudera Runtime 7.3.1 (includes Cloudera Private Cloud Base and Cloudera Public Cloud) for release-specific information related to new features and improvements, bug fixes, deprecated features and components, known issues, and changed features that can affect product behavior.



Attention: Note the following information before proceeding further:

- A new feature introduced in Cloudera Manager 7.13.1 can have a similar impact on the unified Cloudera Runtime 7.3.1 for previous and current Cloudera Manager versions.
- For upgrading Cloudera Manager instructions, see [Upgrading Cloudera Manager 7](#).
- Any changes or modifications made to features in Cloudera Manager 7.13.1 are impacted across unified Cloudera Runtime 7.3.1. For example, a new feature or a configuration change or a behavioral change.
- Any platform support changes made for Cloudera Manager 7.13.1 are impacted across unified Cloudera Runtime 7.3.1. For more information about the supported infrastructure combinations, see [Cloudera support matrix](#).

What's New in Cloudera Manager 7.13.1

Learn about the new features and changed behavior of Cloudera Manager in Cloudera Manager 7.13.1 release.

You must be aware of the additional functionalities and improvements to features of Cloudera Manager in Cloudera Manager 7.13.1. Learn how the new features and improvements benefit you.

New features

Multi Python (Python 3.8 and 3.9) Support for RHEL 8

Cloudera Manager now supports both Python 3.8 and Python 3.9 for RHEL8, providing users with an easy migration path. This support allows users to upgrade to Python 3.9 seamlessly by simply installing Python 3.9 and restarting the Cloudera Manager Agents, with Cloudera Manager automatically detecting and using the highest available Python version.

By maintaining support for both versions, users can upgrade without disrupting cluster operations, ensuring smooth transitions with minimal downtime. This upgrade path helps users stay secure with up-to-date features, security patches, and performance improvements, ensuring their clusters remain stable and future-proof.

For RHEL 8.8 and RHEL 8.10, Cloudera recommends you to install Python 3.9 before upgrading Cloudera Manager to 7.13.1 version to ensure smooth transition with minimal downtime. For information about migrating from Python 3.8 to Python 3.9, see [Migrating from Python 3.8 to Python 3.9 on RHEL 8.8 or RHEL 8.10](#).

cgroup v2 support on RHEL 9 for Cloudera Manager 7.13.1

Cloudera Manager now supports cgroup v2. cgroup v2 offers a unified hierarchy for managing system resources, making it simpler and more efficient compared to cgroup v1. For more information, see [Linux Control Groups \(cgroups\)](#).

You must migrate from cgroup v1 to cgroup v2 for managing the cluster resources using cgroup v2 resource allocation configuration parameters. For information about migrating to cgroup v2, see [Migrating from cgroup v1 to cgroup v2](#).

**Important:**

- Ubuntu 22 is not supported with cgroup v2.
- For the users using RHEL 9.x with Cloudera Manager version lower than 7.13.1, must disable cgroup v2 if already enabled before upgrading to Cloudera Manager 7.13.1 version as cgroup v2 is not supported with Cloudera Manager version lower than 7.13.1.
- During major OS upgrades, while upgrading from Redhat 8 (defaults to cgroup v1) to Redhat 9 (defaults to cgroup v2), the resource configurations will not be automatically transferred such as value of Cgroup V1 CPU Shares will not be populated in Cgroup V2 CPU Weight. Also, the controller files inside the process directories will be created under cgroups root path with default values.
- If you are setting cgroup v1 parameter values manually, then you should now set cgroup v2 parameter values manually (performing conversion of values manually) and restart the services using cgroups.

Note that Cloudera Manager UI will have old values under cgroup v1 parameters which you can use as a reference to re-configure the values in the case of cgroup v2.

Enhancements to the Observability page

The following changes have been made to the Observability page::

- Added role-specific metrics to the Status and Charts Library tabs for component servers such as Pipelines, ADB, and SDX.
- Added relevant metrics across all Observability component servers to the Status and Charts Library tabs for the **Observability** page.

Implemented support for Ranger Plugin Secure Auditing in Solr using Zookeeper.

Support has been added for Ranger plugin secure auditing in Solr by using ZooKeeper.

Added Zookeeper SSL connection support for Ranger & Ranger Raz

Support has been added for ZooKeeper SSL connection for Ranger and Ranger RAZ.

Enhancements to Iceberg replication policies in Replication Manager

The following changes are available for Iceberg replication policies in Replication Manager:

- Added the following options to use during the Iceberg replication policy creation process:
 - JVM Options for Export - You can enter comma-separated JVM options to use for the export process during the Iceberg replication policy run.
 - JVM Options for XFer - You can enter comma-separated JVM options to use for the transfer process during the Iceberg replication policy.
 - JVM Options for Sync - You can enter comma-separated JVM options to use for the sync process during the Iceberg replication policy.
- Iceberg replication policies can replicate V1 and V2 Iceberg tables created using Hive.

What's new in Platform Support

You must be aware of the platform support changes for the Cloudera Manager 7.13.1 release.

This section describes the platform support changes for the Cloudera Manager 7.13.1 associated with Cloudera Private Cloud Base 7.3.1 and Cloudera Public Cloud 7.3.1.

Platform Support Enhancements

- **New OS support:** None
- **New Database support:** None

- **New JDK Version:** None

Fixed Issues in Cloudera Manager 7.13.1

Fixed issues in Cloudera Manager 7.13.1.

OPSAPS-72254: FIPS Failed to upload Spark example jar to HDFS in cluster mode

Fixed an issue with deploying the Spark 3 Client Advanced Configuration Snippet (Safety Valve) for spark3-conf/spark-env.sh.

For more information, see *Added a new Cloudera Manager configuration parameter spark_pyspark_executable_path to Livy for Spark 3* in [Behavioral Changes In Cloudera Manager 7.13.1](#).

OPSAPS-71873 - UCL | CKP4| livyfoo0 kms proxy user is not allowed to access HDFS in 7.3.1.0

In the kms-core.xml file, the Livy proxy user is taken from Livy for Spark 3's configuration in Cloudera 7.3.1 and above.

OPSAPS-70976: The previously hidden real-time monitoring properties are now visible in the Cloudera Manager UI:

The following properties are now visible in the Cloudera Manager UI:

- enable_observability_real_time_jobs
- enable_observability_metrics_dmp

OPSAPS-69996: HBase snapshot creation in Cloudera Manager does not work as expected

During the HBase snapshot creation process, the snapshot create command sometimes tries to create the same snapshot twice because of an unhandled OptimisticLockException during the database write operation. This resulted in intermittent HBase snapshot creation failures. The issue is fixed now.

OPSAPS-66459: Enable concurrent Hive external table replication policies with the same cloud root

When the HIVE_ALLOW_CONCURRENT_REPLICATION_WITH_SAME_CLOUD_ROOT_PATH feature flag is enabled, Replication Manager can run two or more Hive external table replication policies with the same cloud root path concurrently.

For example, if two Hive external table replication policies have s3a://bucket/hive/data as the cloud root path and the feature flag is enabled, Replication manager can run these policies concurrently.

By default, this feature flag is disabled. To enable the feature flag, contact your Cloudera account team.

OPSAPS-70859: Ranger metrics APIs were not working on FedRAMP cluster

On FedRAMP HA cloud cluster, Ranger metrics APIs were not working. This issue is fixed now by introducing new Ranger configurations.

This issue is fixed now by introducing new Ranger configurations.

OPSAPS-71436: Telemetry publisher test Altus connection fails

An error occurred while running the test Altus connection action for Telemetry Publisher. This issue is fixed now.

OPSAPS-68252: The Ranger RMS Database Full Sync command is not visible on cloud clusters

The Ranger RMS Database Full Sync command was not visible on any cloud cluster. Also, it was needed to investigate the minimum user privilege required to see the Ranger RMS Database Full Sync command on the UI.

The issue is fixed now. The command definition on service level in Ranger RMS has been updated after which the command is visible on the UI. The minimum user privilege required to see this command is EnvironmentAdmin.

OPSAPS-69692, OPSAPS-69693: Included filters for Ozone incremental replication in API endpoint

You can use the include filters in the POST `/clusters/{clusterName}/services/{serviceName}/replications` API to replicate only the filtered part of the Ozone bucket. You can use multiple path regular expressions to limit the data to be replicated for an Ozone bucket. For example, if you include the `/path/to/data/.*` and `*/data` filters in the `includeFilter` field for the POST endpoint, the Ozone replication policy replicates only the keys that start with `/path/to/data/.*` or ends with `*/data` in the Ozone bucket.

OPSAPS-70561: Improved page load performance of the “Bucket Browser” tab.

The Cloudera Manager Clusters `[***OZONE SERVICE***]` Bucket Browser tab does not load all the entries of the bucket. Therefore, the page loads faster when you try to display the content of a large bucket with several keys in it.

OPSAPS-71090: The `spark.*.access.hadoopFileSystems` gateway properties are not propagated to Livy.

Added new properties for configuring Spark 2 (`spark.yarn.access.hadoopFileSystems`) and Spark 3 (`spark.kerberos.access.hadoopFileSystems`) that propagate to Livy.

OPSAPS-71271: The `precopylistingcheck` script for Ozone replication policies uses the Ozone replication safety valve value.

The "Run Pre-Filelisting Check" step during Ozone replication uses the content of the `ozone_replication_core_site_safety_valve` property value to configure the Ozone client for the source and the target Cloudera Manager.

OPSAPS-70983: Hive replication command for Sentry to Ranger replication works as expected

The Sentry to Ranger migration during the Hive replication policy run from CDH 6.3.x or higher to CDP Public Cloud 7.3.0.1 or higher is successful.

OPSAPS-69806: Collection of YARN diagnostic bundle will fail

For any combinations of CM 7.11.3 version up to CM 7.11.3 CHF7 version, with CDP 7.1.7 through CDP 7.1.8, collection of the YARN diagnostic bundle will fail, and no data transmits occur.

Now the changes are made to Cloudera Manager to allow the collection of the YARN diagnostic bundle and make this operation successful.

OPSAPS-70655: The `hadoop-metrics2.properties` file is not getting generated into the `ranger-rms-conf` folder

The `hadoop-metrics2.properties` file was getting created in the process directory conf folder, for example, `conf/hadoop-metrics2.properties`, whereas the directory structure in Ranger RMS should be `{process_directory}/ranger-rms-conf/hadoop-metrics2.properties`.

The issue is fixed now. The directory name is changed from `conf` to `ranger-rms-conf`, so that the `hadoop-metrics2.properties` file gets created under the correct directory structure.

OPSAPS-71014: Auto action email content generation failed for some cluster(s) while loading the template file

The issue has been fixed by using a more appropriate template loader class in the freemarker configuration.

OPSAPS-70826: Ranger replication policies fail when target cluster uses Dell EMC Isilon storage and supports JDK17

Ranger replication policies no longer fail if the target cluster is deployed with Dell EMC Isilon storage and also supports JDK17.

OPSAPS-70861: HDFS replication policy creation process fails for Isilon source clusters

When you choose a source CDP Private Cloud Base cluster using the Isilon service and a target cloud storage bucket for an HDFS replication policy in CDP Private Cloud Base Replication Manager UI, the replication policy creation process fails. This issue is fixed now.

OPSAPS-70708: Cloudera Manager Agent not skipping autofs filesystems during filesystem check

Clusters in which there are a large number of network mounts on each host (for example, more than 100 networked file system mounts), cause the startup of Cloudera Manager Agent to take a long time, on the order of 10 to 20 seconds per mount point. This is due to the OS kernel on the cluster host interrogating each network mount on behalf of the Cloudera Manager Agent to gather monitoring information such as file system usage.

This issue is fixed now by adding the ability in the Cloudera Manager Agent's config.ini file to disable filesystem checks.

OPSAPS-68991: Change default SAML response binding to HTTP-POST

The default SAML response binding is HTTP-Artifact, rather than HTTP-POST. While HTTP-POST is designed for handling responses through the POST method, where as HTTP-Artifact necessitates a direct connection with the SP (Cloudera Manager in this case) and Identity Provider (IDP) and is rarely used. HTTP-POST should be the default choice instead.

This issue is fixed now by setting up the new Default SAML Binding to HTTP-POST.

OPSAPS-40169: Audits page does not list failed login attempts on applying Allowed = false filter

The Audits page in Cloudera Manager shows failed login attempts when no filter is applied. However, when the Allowed = false filter is applied it returns 0 results. Whereas it should have listed those failed login attempts. This issue is fixed now.

OPSAPS-70583: File Descriptor leak from Cloudera Manager 7.11.3 CHF3 version to Cloudera Manager 7.11.3 CHF7

Unable to create NettyTransceiver due to Avro library upgrade which leads to File Descriptor leak. File Descriptor leak occurs in Cloudera Manager when a service tries to talk with Event Server over Avro. This issue is fixed now.

OPSAPS-70962: Creating a cloud restore HDFS replication policy with a peer cluster as destination which is not supported by Replication Manager

During the HDFS replication policy creation process, incorrect Destination clusters and MapReduce services appear which when chosen creates a dummy replication policy to replicate from a cloud account to a remote peer cluster. This scenario is not supported by Replication Manager. This issue is now fixed.

OPSAPS-71108: Use the earlier format of PCR

You can use the latest version of the PCR (Post Copy Reconciliation) script, or you can restore PCR to the earlier format by setting the `com.cloudera.enterprise.distcp.post-copy-reconciliation.legacy-output-format.enabled=true` key value pair in the Cloudera Manager Clusters *HDFS service* Configuration `hdfs_replication_hdfs_site_safety_valve` property.

OPSAPS-70689: Enhanced performance of DistCp CRC check operation

When a MapReduce job for an HDFS replication policy job fails, or when there are target-side changes during a replication job, Replication Manager initiates the bootstrap replication process. During this process, a cyclic redundancy check (CRC) check is performed by default to determine whether a file can be skipped for replication.

By default, the CRC for each file is queried by the mapper (running on the target cluster) from the source cluster's NameNode. The round trip between the source and target cluster for each file consumes network resources and raises the cost of execution. To improve the performance, you can set the following variables to true, on the target cluster, to improve the performance of the CRC check for the Cloudera Manager Clusters *HDFS service* Configuration `HDFS_REPLICATION_ENV_SAFETY_VALVE` property:

- `ENABLE_FILESTATUS_EXTENSIONS`
- `ENABLE_FILESTATUS_CRC_EXTENSIONS`

By default, these are set to false.

After you set the key-value pairs, the CRC for each file is queried locally from the NameNode on the source cluster and copied over to the target cluster at the end of the replication process, which reduces the cost because round trip is between two nodes of the same cluster. The CRC checksums are written to the file listing files.

OPSAPS-70685: Post Copy Reconciliation (PCR) for HDFS replication policies between on-premises clusters

To add the Post Copy Reconciliation (PCR) script to run as a command step during the HDFS replication policy job run, you can enter the `SCHEDULES_WITH_ADDITIONAL_DEBUG_STEPS = [***ENTER COMMA-SEPARATED LIST OF NUMERICAL IDS OF THE REPLICATION POLICIES***]` key-value pair in the target Cloudera Manager Clusters *HDFS service* `hdfs_replication_env_safety_valve` property.

To run the PCR script on the HDFS replication policy, use the `/clusters/[***CLUSTER NAME***]/>/services/[***SERVICE***]/replications/[***SCHEDULE ID***]/postCopyReconciliation` API.

For more information about the PCR script, see [How to use the post copy reconciliation script for HDFS replication policies](#).

OPSAPS-70188: Conflicts field missing in ParcelInfo

Fixed an issue in parcels where conflicts field in manifest.json would mark a parcel as invalid

OPSAPS-70248: Optimize Impala Graceful Shutdown Initiation Time

This issue is resolved by streamlining the shutdown initiation process, reducing delays on large clusters.

OPSAPS-70157: Long-term credential-based GCS replication policies continue to work when cluster-wide IDBroker client configurations are deployed

Replication policies that use long-term GCS credentials work as expected even when cluster-wide IDBroker client configurations are configured.

OPSAPS-70422: Change the “Run as username(on source)” field during Hive external table replication policy creation

You can use a different user other than `hdfs` for Hive external table replication policy run to replicate from an on-premises cluster to the cloud bucket if the `USE_PROXY_USER_FOR_CLOUD_TRANSFER=true` key-value pair is set for the source Cloudera Manager Clusters *Hive service* Configuration Hive Replication Environment Advanced Configuration Snippet (Safety Valve) property. This is applicable for all external accounts other than IDBroker external account.

OPSAPS-70460: Allow white space characters in Ozone snapshot-diff parsing

Ozone incremental replication no longer fails if a changed path contains one or more space characters.

OPSAPS-70594: Ozone HttpFS gateway role is not added to Rolling Restart

This issue is now resolved by adding the Ozone HttpFS gateway role to the Rolling Restart.

OPSAPS-68752: Snapshot-diff delta is incorrectly renamed/deleted twice during on-premises to cloud replication

The snapshots created during replication are deleted twice instead of once, which results in incorrect snapshot information. This issue is fixed. For more information, see [Cloudera Customer Advisory 2023-715: Replication Manager may delete its snapshot information when migrating from on-prem to cloud](#).

OPSAPS-70226: Atlas uses the Solr configuration directory available in ATLAS_PROCESS/conf/solr instead of the Cloudera Manager provided directory

Atlas uses the configuration in `/var/run/cloudera-scm-agent/process/151-atlas-ATLAS_SERVER/solrconf.xml`.

OPSAPS-68112: Atlas diagnostic bundle should contain server log, configurations, and, if possible, heap memories

The diagnostic bundle contains server log, configurations, and heap memories in a GZ file inside the diagnostic .zip package.

OPSAPS-69921: ATLAS_OPTS environment variable is set for FIPS with JDK 11 environments to run the import script in Atlas

_JAVA_OPTIONS are populated with additional parameters as seen in the following:

```
java_opts = 'export _JAVA_OPTIONS="-Dcom.safelogic.cryptocomply.fips.approved_only=true ' \
'--add-modules=com.safelogic.cryptocomply.fips.core,' \
'bctls --add-exports=java.base/sun.security.provider=com.safelogic.cryptocomply.fips.core ' \
'--add-exports=java.base/sun.security.provider=bctls --module-path=/cdep/extra_jars ' \
'-Dcom.safelogic.cryptocomply.fips.approved_only=true -Djdk.tls.ephemeralDHKeySize=2048 ' \
'-Dorg.bouncycastle.jsse.client.assumeOriginalHostName=true -Djdk.tls.trustNameService=true" '
```

OPSAPS-71258: Kafka, SRM, and SMM cannot process messages compressed with Zstd or Snappy if /tmp is mounted as noexec

The issue is fixed by using JVM flags that point to a different temporary folder for extracting the native library.

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

Cloudera Manager now registers the metrics kafka_connect_connector_task_metrics_batch_size_avg and kafka_connect_connector_task_metrics_batch_size_max correctly.

OPSAPS-68708: Schema Registry might fail to start if a load balancer address is specified in Ranger

Schema Registry now always ensures that the address it uses to connect to Ranger ends with a trailing slash (/). As a result, Schema Registry no longer fails to start if Ranger has a load balancer address configured that does not end with a trailing slash.

OPSAPS-69978: Cruise Control capacity.py script fails on Python 3

The script querying the capacity information is now fully compatible with Python 3.

OPSAPS-64385: Atlas's client.auth.enabled configuration is not configurable

In customer environments where user certifications are required to authenticate to services, the Apache Atlas web UI will constantly prompt for certifications. To solve this, the client.auth.enabled parameter is set to true by default. If it is needed to set it false, then you need to override the setting from safety-valve with a configuration snippet. Once it set to false, then no more certificate prompts will be displayed.

OPSAPS-71089: Atlas's client.auth.enabled configuration is not configurable

In customer environments where user certifications are required to authenticate to services, the Apache Atlas web UI will constantly prompt for certifications. To solve this, the client.auth.enabled parameter is set to true by default. If it is needed to set it false, then you need to override the setting from safety-valve with a configuration snippet. Once it set to false, then no more certificate prompts will be displayed.

OPSAPS-71677: When you are upgrading from CDP Private Cloud Base 7.1.9 SP1 to CDP Private Cloud Base 7.3.1, upgrade-rollback execution fails during HDFS rollback due to missing directory.

This issue is now resolved. The HDFS meta upgrade command is executed by creating the previous directory due to which the rollback does not fail.

OPSAPS-71390: COD cluster creation is failing on INT and displays the Failed to create HDFS directory /tmp error.

This issue is now resolved. Export options for jdk17 is added.

OPSAPS-71188: Modify default value of dfs_image_transfer_bandwidthPerSec from 0 to a feasible value to mitigate RPC latency in the namenode.

This issue is now resolved.

OPSAPS-58777: HDFS Directories are created with root as user.

This issue is now resolved by fixing service.sdl.

OPSAPS-71474: In Cloudera Manager UI, the Ozone service Snapshot tab displays label label.goToBucket and it must be changed to Go to bucket.

This issue is now resolved.

OPSAPS-70288: Improvements in master node decommissioning.

This issue is now resolved by making usability and functional improvements to the Ozone master node decommissioning.

OPSAPS-71647: Ozone replication fails for incompatible source and target Cloudera Manager versions during the payload serialization operation

Replication Manager now recognizes and annotates the required fields during the payload serialization operation. For the list of unsupported Cloudera Manager versions that do not have this fix, see [Preparing clusters to replicate Ozone data](#).

OPSAPS-71156: PostCopyReconciliation ignores mismatching modification time for directories

The Post Copy Reconciliation script (PCR) script does not check the file length, last modified time, and cyclic redundancy check (CRC) checksums for directories (paths that are directories) on both the source and target clusters.

OPSAPS-70732: Atlas replication policies no longer consider inactive Atlas server instances

Replication Manager considers only the active Atlas server instances during Atlas replication policy runs.

OPSAPS-70924: Configure Iceberg replication policy level JVM options

You can add replication-policy level JVM options for the export, transfer, and sync CLIs for Iceberg replication policies on the **Advanced** tab in the **Create Iceberg Replication Policy** wizard.

OPSAPS-70657: KEYTRUSTEE_SERVER & RANGER_KMS_KTS migration to RANGER_KMS from CDP 7.1.x to UCL

KEYTRUSTEE_SERVER and RANGER_KMS_KTS services are not supported starting from the CDP 7.3.1 release. Therefore added validation and confirmation messages to the CM upgrade wizard to alert the user to migrate KEYTRUSTEE_SERVER keys to RANGER_KMS before upgrading to CDP 7.3.1 release.

OPSAPS-70656: Remove KEYTRUSTEE_SERVER & RANGER_KMS_KTS from CM for UCL

The Keytrustee components - KEYTRUSTEE_SERVER and RANGER_KMS_KTS services are not supported starting from the CDP 7.3.1 release. These services cannot be installed or managed with CM 7.13.1 using CDP 7.3.1.

OPSAPS-67480: In 7.1.9, default Ranger policy is added from the cdp-proxy-token topology, so that after a new installation of CDP-7.1.9, the Knox-Ranger policy includes cdp-proxy-token. However, upgrades do not add cdp-proxy-token to cm_knox policies automatically.

This issue is fixed now.

OPSAPS-70838: Flink user should be add by default in ATLAS_HOOK topic policy in Ranger >> cm_kafka

The "flink" service user is granted publish access on the ATLAS_HOOK topic by default in the Kafka Ranger policy configuration.

OPSAPS-69411: 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 and tables and the respective URLs) by using the config "authorization.migration.export.migration_objects" during export.

OPSAPS-68252: "Ranger RMS Database Full Sync" option was not visible on mow-int cluster setup for hrt_qa user (7.13.0.0)

The fix makes the command visible on cloud clusters when the user has minimum EnvironmentAdmin privilege.

OPSAPS-70148: Ranger audit collection creation is failing on latest SSL enabled UCL cluster due to zookeeper connection issue

Added support for secure ZooKeeper connection for the Ranger Plugin Solr audit connection configuration xasecure.audit.destination.solr.zookeepers.

OPSAPS-52428: Add SSL to ZooKeeper in CDP

Added SSL/TLS encryption support to CDP components. ZooKeeper SSL (secure) port now gets automatically enabled and components communicate on the encrypted channel if cluster has AutoTLS enabled.

OPSAPS-72093: FIPS - yarn jobs are failing with No key provider is configured

The yarn.nodemanager.admin environment must contain the FIPS related Java options, and this configuration is handled such that the comma is a specific character in the string. This change proposes to use single module additions in the default FIPS options (use separate --add-modules for every module), and it adds the FIPS options to the yarn.nodemanager.admin environment.

Previously, yarn.nodemanager.container-localizer.admin.java.opts contained FIPS options only for 7.1.9, this patch also fixes this, and adds the proper configurations in 7.3.1 environments also.

This was tested on a real cluster, and with the current changes YARN works properly, and can successfully run distcp from/to encryption zones.

OPSAPS-70113: Fix the ordering of YARN admin ACL config

The YARN Admin ACL configuration in Cloudera Manager shuffled the ordering when it was generated. This issue is now fixed, so that the input ordering is maintained and correctly generated.

Known Issues in Cloudera Manager 7.13.1

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Manager 7.13.1.

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.

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_transport` in the configuration file located at `/etc/my.cnf`.

OPSAPS-70771: Running Ozone replication policy does not show performance reports

During an Ozone replication policy run, the A server error has occurred. See Cloudera Manager server log for details error message appears when you click:

- Performance Reports OZONE Performance Summary or Performance Reports OZONE Performance Full on the **Replication Policies** page.
- **Download CSV** on the **Replication History** page to download any report.

None

CDPD-53185: Clear REPL_TXN_MAP table on target cluster when deleting a Hive ACID replication policy

The entry in `REPL_TXN_MAP` table on the target cluster is retained when the following conditions are true:

1. A Hive ACID replication policy is replicating a transaction that requires multiple replication cycles to complete.
2. The replication policy and databases used in it get deleted on the source and target cluster even before the transaction is completely replicated.

In this scenario, if you create a database using the same name as the deleted database on the source cluster, and then use the same name for the new Hive ACID replication policy to replicate the database, the replicated database on the target cluster is tagged as 'database incompatible'. This happens after the housekeeper thread process (that runs every 11 days for an entry) deletes the retained entry.

Create another Hive ACID replication policy with a different name for the new database

OPSAPS-71592: Replication Manager does not read the default value of “ozone_replication_core_site_safety_valve” during Ozone replication policy run

During the Ozone replication policy run, Replication Manager does not read the value in the `ozone_replication_core_site_safety_valve` advanced configuration snippet if it is configured with the default value.

To mitigate this issue, you can use one of the following methods:

- Remove some or all the properties in `ozone_replication_core_site_safety_valve`, and move them to `ozone-conf/ozone-site.xml_service_safety_valve`.
- Add a dummy property with no value in `ozone_replication_core_site_safety_valve`. For example, add `<property><name>dummy_property</name><value></value></property>`, save the changes, and run the Ozone replication policy.

OPSAPS-71897: Finalize Upgrade command fails after upgrading the cluster with CustomKerberos setup causing INTERNAL_ERROR with EC writes.

After the UI `FinalizeCommand` fails, you must manually run the `finalize` commands through the Ozone Admin CLI:

1. `kinit` with the `scm custom kerberos principal`

2. `ozone admin scm finalizeupgrade`
3. `ozone admin scm finalizationstatus`

OPSAPS-70702: Ranger replication policies fail because of the truststore file location

Ranger replication policies fail during the Exporting services, policies and roles from Ranger remote step.

- Log in to the Ranger Admin host(s) on the source cluster.
- Identify the Cloudera Manager agent PEM file using the `# cat /etc/cloudera-scm-agent/config.ini | grep -i client_cert_file` command. For example, the file might reside in `client_cert_file=/myTLSpath/cm_server-cert.pem` location.
- Create the path for the new PEM file using the `# mkdir -p /var/lib/cloudera-scm-agent/agent-cert/` command.
- Copy the `client_cert_file` from `config.ini` as `cm-auto-global_cacerts.pem` file using the `# cp /myTLSpath/cm_server-cert.pem /var/lib/cloudera-scm-agent/agent-cert/cm-auto-global_cacerts.pem` command.
- Change the ownership to 644 using the `# chmod 644 /var/lib/cloudera-scm-agent/agent-cert/cm-auto-global_cacerts.pem` command.
- Resume the Ranger replication policy in Replication Manager.



Note: Ensure that you change `/myTLSpath/cm_server-cert.pem` to the actual PEM file location defined in `config.ini` under `client_cert_file`.

OPSAPS-71424: The configuration sanity check step ignores during the replication advanced configuration snippet values during the Ozone replication policy job run

The OBS-to-OBS Ozone replication policy jobs fail if the S3 property values for `fs.s3a.endpoint`, `fs.s3a.secret.key`, and `fs.s3a.access.key` are empty in Ozone Service Advanced Configuration Snippet (Safety Valve) for `ozone-conf/ozone-site.xml` even though you defined the properties in Ozone Replication Advanced Configuration Snippet (Safety Valve) for `core-site.xml`.

Ensure that the S3 property values for `fs.s3a.endpoint`, `fs.s3a.secret.key`, and `fs.s3a.access.key` contains at least a dummy value in Ozone Service Advanced Configuration Snippet (Safety Valve) for `ozone-conf/ozone-site.xml`.

Additionally, you must ensure that you do not update the property values in Ozone Replication Advanced Configuration Snippet (Safety Valve) for `core-site.xml` for Ozone replication jobs. This is because the values in this advanced configuration snippet overrides the property values in `core-site.xml` and not the `ozone-site.xml` file.

Different property values in Ozone Service Advanced Configuration Snippet (Safety Valve) for `ozone-conf/ozone-site.xml` and Ozone Replication Advanced Configuration Snippet (Safety Valve) for `core-site.xml` result in a nondeterministic behavior where the replication job picks up either value during the job run which leads to incorrect results or replication job failure.

OPSAPS-71403: Ozone replication policy creation wizard shows "Listing Type" field in source Cloudera Private Cloud Base versions lower than 7.1.9

When the source Cloudera Private Cloud Base cluster version is lower than 7.1.9 and the Cloudera Manager version is 7.11.3, the Ozone replication policy creation wizard shows Listing Type and its options. These options are not available in Cloudera Private Cloud Base 7.1.8x versions.

OPSAPS-71659: Ranger replication policy fails because of incorrect source to destination service name mapping

Ranger replication policy fails because of incorrect source to destination service name mapping format during the transform step.

If the service names are different in the source and target, then you can perform the following steps to resolve the issue:

1. SSH to the host on which the Ranger Admin role is running.

2. Find the ranger-replication.sh file.
3. Create a backup copy of the file.
4. Locate substituteEnv SOURCE_DESTINATION_RANGER_SERVICE_NAME_MAPPING \${RANGER_REPL_SERVICE_NAME_MAPPING} in the file.
5. Modify it to substituteEnv
SOURCE_DESTINATION_RANGER_SERVICE_NAME_MAPPING
"\${RANGER_REPL_SERVICE_NAME_MAPPING/\/}"
6. Save the file.
7. Rerun the Ranger replication policy.

OPSAPS-69782: HBase COD-COD replication from 7.3.1 to 7.2.18 fails during the "create adhoc snapshot" step

An HBase replication policy replicating from 7.3.1 COD to 7.2.18 COD cluster that has ‘Perform Initial Snapshot’ enabled fails during the snapshot creation step in Cloudera Replication Manager.

OPSAPS-71414: Permission denied for Ozone replication policy jobs if the source and target bucket names are identical

The OBS-to-OBS Ozone replication policy job fails with the com.amazonaws.services.s3.model.AmazonS3Exception: Forbidden or Permission denied error when the bucket names on the source and target clusters are identical and the job uses S3 delegation tokens. Note that the Ozone replication jobs use the delegation tokens when the S3 connector service is enabled in the cluster.

You can use one of the following workarounds to mitigate the issue:

- Use different bucket names on the source and target clusters.
- Set the fs.s3a.delegation.token.binding property to an empty value in ozone_replication_core_site_safety_valve to disable the delegation tokens for Ozone replication policy jobs.

OPSAPS-71256: The “Create Ranger replication policy” action shows 'TypeError' if no peer exists

When you click target Cloudera Manager Replication Manager Replication Policies Create Replication Policy Ranger replication policy , the TypeError: Cannot read properties of undefined error appears.

OPSAPS-71067: Wrong interval sent from the Replication Manager UI after Ozone replication policy submit or edit process.

When you edit the existing Ozone replication policies, the schedule frequency changes unexpectedly.

OPSAPS-70848: Hive external table replication policies fail if the source cluster is using Dell EMC Isilon storage

During the Hive external table replication policy run, the replication policy fails at the Hive Replication Export step. This issue is resolved.

OPSAPS-71005: RemoteCmdWork uses a singlethreaded executor

Replication Manager runs the remote commands for a replication policy through a single-thread executor.

OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

Workaround: SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You need to override the bootstrap server URL by performing the following steps:

1. In Cloudera Manager, go to SMM Configuration Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve)
2. Override bootstrap server URL (hostname:port as set in the listeners for broker) for streams-messaging-manager.yaml.

3. Save your changes.
4. Restart SMM.

OPSAPS-69317: Kafka Connect Rolling Restart Check fails if SSL Client authentication is required

The rolling restart action does not work in Kafka Connect when the `ssl.client.auth` option is set to required. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set `ssl.client.auth` to requested instead of required and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

OPSAPS-70971: Schema Registry does not have permissions to use Atlas after an upgrade

Following an upgrade, Schema Registry might not have the required permissions in Ranger to access Atlas. As a result, Schema Registry's integration with Atlas might not function in secure clusters where Ranger authorization is enabled.

1. Access the Ranger Console (Ranger Admin web UI).
2. Click the `cm_atlas` resource-based service.
3. Add the `schemaregistry` user to the all - * policies.
4. Click Manage Service Edit Service .
5. Add the `schemaregistry` user to the `default.policy.users` property.

OPSAPS-59597: SMM UI logs are not supported by Cloudera Manager

Cloudera Manager does not display a Log Files menu for SMM UI role (and SMM UI logs cannot be displayed in the Cloudera Manager UI) because the logging type used by SMM UI is not supported by Cloudera Manager.

View the SMM UI logs on the host.

OPSAPS-72298: Impala metadata replication is mandatory and UDF functions parameters are not mapped to the destination

Impala metadata replication is enabled by default but the legacy Impala C/C++ UDF's (user-defined functions) are not replicated as expected during the Hive external table replication policy run.

Edit the location of the UDF functions after the replication run is complete. To accomplish this task, you can edit the "path of the UDF function" to map it to the new cluster address, or you can use a script.

OPSAPS-70713: Error appears when running Atlas replication policy if source or target clusters use Dell EMC Isilon storage

You cannot create an Atlas replication policy between clusters if one or both the clusters use Dell EMC Isilon storage.

None

OPSAPS-72468: Subsequent Ozone OBS-to-OBS replication policy do not skip replicated files during replication

The first Ozone replication policy run is a bootstrap run. Sometimes, the subsequent runs might also be bootstrap jobs if the incremental replication fails and the job runs fall back to bootstrap replication. In this scenario, the bootstrap replication jobs might replicate the files that were already replicated because the modification time is different for a file on the source and the target cluster.

None

OPSAPS-72470: Hive ACID replication policies fail when target cluster uses Dell EMC Isilon storage and supports JDK17

Hive ACID replication policies fail if the target cluster is deployed with Dell EMC Isilon storage and also supports JDK17.

None

Behavioral Changes In Cloudera Manager 7.13.1

You can review the changes in certain features or functionalities of Cloudera Manager that have resulted in a change in behavior from the previously released version to this version of Cloudera Manager 7.13.1.

Added ability in the Cloudera Manager Agent's config.ini file to disable filesystem checks.

In Cloudera Manager Agent 7.13.1 and higher versions, a new optional configuration flag is available. The new flag is `monitor_filesystems`, which you can set up in the Cloudera Manager Agent `config.ini` file (found in `/etc/cloudera-scm-agent/config.ini`).

You can add the following lines in the `config.ini` file before upgrading Cloudera Manager Agent to disable monitoring of filesystems:

- The flag `monitor_filesystems` is used to determine if the agent has to monitor the filesystems.
- If the flag is set to `True`, Cloudera Manager Agent monitors the filesystems.
- If the flag is set to `False`, Cloudera Manager Agent will not monitor any filesystems. If the flag is not included in the file, it will default to `True`, and Cloudera Manager Agent behavior will match previous versions.



Attention: The side-effect of this change is that Cloudera Manager Server will not display filesystem usage for any filesystem (local or networked) for the modified host. A future version of Cloudera Manager Agent will have changes to specifically avoid networked filesystems, while still monitoring local filesystems.

Added a new Cloudera Manager configuration parameter `spark_pyspark_executable_path` to Livy for Spark 3.

In Cloudera Manager Agent 7.13.1 and higher versions, a new Cloudera Manager configuration parameter `spark_pyspark_executable_path` is added to Livy for Spark 3 service.

The value of `spark_pyspark_executable_path` for Livy must sync with the value of the Spark 3 service's `spark_pyspark_executable_path` parameter in Cloudera Manager.



Important:

If the `PYSPARK_PYTHON/PYSPARK_DRIVER_PYTHON` environment variables are not set in `spark-env.sh`, then the default value of these variables will be the value of the `spark_pyspark_executable_path` Cloudera Manager property.

The default value of `spark_pyspark_executable_path` is `/opt/cloudera/cm-agent/bin/python`.

Summary: The Livy proxy user is taken from Livy for Spark 3's configuration.

Previous behavior:

The custom Kerberos principal configuration was updated via the Livy service.

New behavior:

The Livy proxy user is taken from Livy for Spark 3's configuration, as the Livy service has been replaced with Livy for Spark3 in Cloudera Private Cloud Public Cloud version 7.3.1.

Fixed Common Vulnerabilities and Exposures in Cloudera Manager 7.13.1

Common Vulnerabilities and Exposures (CVE) that are fixed in Cloudera Manager 7.13.1 associated with Cloudera Private Cloud Base 7.3.1 and Cloudera Public Cloud 7.3.1.

Cloudera Manager 7.13.1

CVEs	Package Name
CVE-2019-14893	Jackson-databind

CVEs	Package Name
CVE-2020-9546	Jackson-databind
CVE-2020-10672	Jackson-databind
CVE-2020-10968	Jackson-databind
CVE-2020-10969	Jackson-databind
CVE-2020-11111	Jackson-databind
CVE-2020-11112	Jackson-databind
CVE-2020-11113	Jackson-databind
CVE-2020-11619	Jackson-databind
CVE-2020-11620	Jackson-databind
CVE-2020-14060	Jackson-databind
CVE-2020-14061	Jackson-databind
CVE-2020-14062	Jackson-databind
CVE-2020-14195	Jackson-databind
CVE-2020-35728	Jackson-databind
CVE-2020-25649	Jackson-databind
CVE-2021-29425	commons-io
CVE-2021-28168	Jersey
CVE-2023-33202	Bouncycastle
CVE-2024-34447	Bouncycastle
CVE-2024-29857	Bouncycastle
CVE-2024-30171	Bouncycastle
CVE-2023-33201	Bouncycastle
CVE-2020-11971	Apache Camel
CVE-2018-1282	Apache Hive
CVE-2018-11777	Apache Hive
CVE-2021-34538	Apache Hive
CVE-2020-1926	Apache Hive
CVE-2018-1314	Apache Hive
CVE-2018-1284	Apache Hive
CVE-2018-1315	Apache Hive
CVE-2021-46877	Jackson-databind
CVE-2020-13697	Nanohttpd
CVE-2022-21230	Nanohttpd
CVE-2024-29736	Apache CXF
CVE-2024-32007	Apache CXF
CVE-2022-1415	Drools
CVE-2021-41411	Drools
CVE-2018-8009	Apache Hadoop
CVE-2014-3577	Apache httpclient

CVEs	Package Name
CVE-2015-5262	Apache httpclient
CVE-2016-6811	Apache Hadoop
CVE-2018-8029	Apache Hadoop
CVE-2018-11768	Apache Hadoop
CVE-2018-1296	Apache Hadoop
CVE-2017-3162	Apache Hadoop
CVE-2017-15713	Apache Hadoop
CVE-2017-3161	Apache Hadoop
CVE-2016-5001	Apache Hadoop
CVE-2016-3086	Apache Hadoop
CVE-2016-5393	Apache Hadoop
CVE-2024-23454	Apache Hadoop
CVE-2018-11765	Apache Hadoop
CVE-2020-9492	Apache Hadoop
CVE-2015-1776	Apache Hadoop
CVE-2016-10735	Bootstrap
CVE-2018-14041	Bootstrap
CVE-2018-14042	Bootstrap
CVE-2018-20676	Bootstrap
CVE-2018-20677	Bootstrap
CVE-2019-8331	Bootstrap
CVE-2020-28458	Datatables
CVE-2021-23445	Datatables
CVE-2015-6584	Datatables
CVE-2016-4055	moment.js
CVE-2019-20444	Netty
CVE-2019-20445	Netty
CVE-2015-2156	Netty
CVE-2016-4970	Netty
CVE-2019-16869	Netty
CVE-2020-7238	Netty
CVE-2021-37136	Netty
CVE-2021-37137	Netty
CVE-2022-41881	Netty
CVE-2021-43797	Netty
CVE-2023-34462	Netty
CVE-2021-21295	Netty
CVE-2021-21409	Netty
CVE-2021-21290	Netty

CVEs	Package Name
CVE-2022-24823	Netty
CVE-2017-3166	Apache Hadoop
CVE-2017-15718	Apache Hadoop
CVE-2018-8025	Apache Hbase
CVE-2019-0212	Apache Hbase
CVE-2022-25647	Gson
CVE-2019-9518	Netty
CVE-2020-11612	Netty
CVE-2016-5724	Cloudera CDH
CVE-2017-9325	Cloudera CDH
CVE-2021-41561	Apache Parquet
CVE-2022-26612	Apache Hadoop
CVE-2024-36124	Snappy
CVE-2015-7521	Apache Hive
CVE-2016-3083	Apache Hive
CVE-2015-1772	Apache Hive
CVE-2022-41853	hsqldb
CVE-2015-8094	Cloudera Hue
CVE-2021-28170	javax.el
CVE-2011-4461	Mortbay Jetty
CVE-2009-1523	Mortbay Jetty
CVE-2023-5072	org.json
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-2023-0833	Okhttp
CVE-2023-52428	Nimbus-jose-jwt
CVE-2021-0341	Okhttp
CVE-2018-11799	Apache Oozie
CVE-2017-15712	Apache Oozie
CVE-2024-1597	Postgresql
CVE-2022-34169	Apache Xalan
CVE-2022-1471	Snakeyaml
CVE-2023-43642	Snappy Java
CVE-2022-22965	Spring Framework

CVEs	Package Name
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
CVE-2023-20862	Spring Security
CVE-2024-22257	Spring Security
CVE-2023-20859	Spring Vault
CVE-2024-22243	Spring Framework
CVE-2024-22262	Spring Framework
CVE-2023-44981	Apache Zookeeper
CVE-2016-5017	Apache Zookeeper
CVE-2018-8012	Apache Zookeeper
CVE-2019-0201	Apache Zookeeper

Deprecation notices in Cloudera Manager 7.13.1

Certain features and functionalities have been removed or deprecated in Cloudera Manager 7.13.1. 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.

Platform and OS

The listed Operating Systems and databases are deprecated or removed from the Cloudera Manager 7.13.1 release.

Database Support

The following databases are removed and no longer supported from the Cloudera Manager 7.13.1 release:

- PostgreSQL 12
- MariaDB 10.4
- MySQL 5.7

Operating System

The following operating systems are removed and no longer supported from the Cloudera Manager 7.13.1 release:

- RHEL 8.6
- RHEL 7.9
- RHEL 7.9 (FIPS)
- CentOS 7.9
- SLES 12 SP5