Cloudera Manager 7.13.1

Unified Cloudera Manager Release Notes

Date published: 2024-12-10
Date modified: 2025-03-18

CLOUDERA

https://docs.cloudera.com/

Legal Notice

Contents

Cloudera Manager 7.13.1 Release Notes

This document provides a summary of the latest updates in Cloudera Manager 7.13.1 and its Cumulative hotfixes. You can review the Release Notes of Cloudera Manager 7.13.1 associated with unified Cloudera Runtime 7.3.1 (includes Cloudera Base on premises and Cloudera on 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.

Important:

- Cloudera Manager 7.13.1 CHF5 (7.13.1.501) supports Cloudera Data Services on premises 1.5.5 SP1 release.
- Cloudera Manager 7.13.1 CHF4 (7.13.1.402) supports Cloudera Data Services on premises 1.5.5 CHF1 release.
- Cloudera Manager 7.13.1 CHF3 (7.13.1.303) supports Cloudera Data Services on premises 1.5.5 release.
- Cloudera Manager 7.13.1 CHF1 (7.13.1.100) supports Cloudera Data Services on premises 1.5.4 SP2 release.
- Cloudera Runtime 7.3.1.400 SP2 does not support running on hosts that contain Python 3.11. To avoid issues, you must ensure that Python 3.11 is not installed on any hosts under management by Cloudera Manager 7.13.1 CHF4 (7.13.1.400).
- Contact Cloudera Support for questions related to any specific hotfixes.

Important:

From Cloudera Manager 7.13.1 onwards Cloudera Manager by default does not support Cloudera Runtime 6 cluster because CDH 6 jars have security vulnerabilities.

If you want to run Cloudera Manager 7.13.1 or higher versions with CDH 6 cluster then you need to install an additional cloudera-manager-daemons-cdh6 rpm/debian package on every cluster host after installation/ upgrade to Cloudera Manager 7.13.1. This package is available along with the other Cloudera Manager packages on the public repository.

Important: Note the following information before proceeding further:

- 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

Learn about the new features and changed behavior of Cloudera Manager in Cloudera Manager 7.13.1 and its cumulative hotfixes.

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

Cloudera Manager 7.13.1 Cumulative hotfix 5 (7.13.1.500)

Python 3.11 Support for Cloudera Runtime 7.3.1.500 SP3

Cloudera Manager 7.13.1 CHF5 now supports Python 3.11 when using Cloudera Runtime 7.3.1.500 SP3 along with Cloudera Runtime 7.1.9 SP1 CHF9 or higher versions.

Cloudera recommends you to install Python 3.9 or Python 3.11 before upgrading Cloudera Manager to 7.13.1 CHF5 version to ensure smooth transition with minimal downtime. For information about migrating Python for various operating systems, see Migrating Python versions.

Added feature flag for Key Version Numbers (KVNOs) validation functionality

Cloudera Manager validates Key Version Numbers (KVNOs) in keytabs against entries fetched from KDC server by storing it in its internal CREDENTIALS database to detect stale or outdated keytabs, improving Kerberos security and reliability. By default, this validation is enabled (

CMF_FF_KERBEROS_KVNO_VALIDATION=true

). For environments using externally managed KDCs or keytabs, this check can be disabled by setting the flag to

false

in

/etc/default/cloudera-scm-server

file. This ensures flexibility for both Cloudera Manager-managed and externally managed Kerberos infrastructures.

Trino plugin support in Ranger

Trino plugin support in Ranger has been added.

Cloudera Manager 7.13.1 Cumulative hotfix 4 (7.13.1.400)

Extended Multi Python (Python 3.8, 3.9, 3.10, and 3.11) Support in Cloudera Manager 7.13.1 CHF4

Cloudera Manager 7.13.1 CHF4 significantly expands its support for multiple Python environments when used in combination with Cloudera Runtime 7.1.9 SP1 CHF9 and higher versions. Previously, multi-Python was limited to RHEL 8, with Python 3.8 as default and an optional upgrade to Python 3.9.

Important:

- Cloudera Manager 7.13.1 CHF4 supports Python 3.11 only when using Cloudera Runtime 7.1.9 SP1 CHF9.
- Do not install Python 3.11 on hosts running Cloudera Runtime 7.3.1.400 SP2.

The specific Python versions supported are dependent on your operating system. The Cloudera Manager Agent automatically selects the highest available Python version from the supported range of Python versions (Python 3.8 - Python 3.11) that is installed on the host after the Cloudera Manager Agent restart. This intelligent selection simplifies management by ensuring the Cloudera Manager Agents always use the most up-to-date compatible Python for the specific Operating System.

The Python selection assumes a Python installation through the OS packages.

Important:

You must ensure that the Python version you intend for Cloudera Manager to use is the highest version installed on the host within the supported range.

Example:

If you are on Cloudera Manager 7.13.1.400 and have installed all Python versions (supported range: Python 3.8 - Python 3.11) on a host:

- The Cloudera Manager Agent will automatically select and use Python 3.11 after a restart, as it is the highest available version.
- If your intention is for the Cloudera Manager Agent to use Python 3.9 instead, you must uninstall Python 3.11 from that host and restart the Cloudera Manager Agent.

The following table provides the details about the Python versions that are supported for various Operating Systems when Cloudera Manager 7.13.1 CHF4 used in combination with various Cloudera Runtime versions:

Operating system	Supported Python versions for (Cloudera Manager 7.13.1 CHF4 + Cloudera Runtime 7.1.9 SP1 CHF9) combination	Supported Python versions for (Cloudera Manager 7.13.1 CHF4 + Cloudera Runtime 7.3.1.400 SP2) combination
RHEL 8	• Python 3.8 • Python 3.9 • Python 3.11	• Python 3.8 • Python 3.9
RHEL 9	• Python 3.9 • Python 3.11	Python 3.9 only
Oracle Linux 8	• Python 3.8 • Python 3.9 • Python 3.11	Python 3.8 only
Oracle Linux 9	• Python 3.9 • Python 3.11	Python 3.9 only
Rocky Linux 9	• Python 3.9 • Python 3.11	Python 3.9 only
SLES 15 SP4	Python 3.11 only	Python 3.10 only
SLES 15 SP5	Python 3.11 only	Python 3.10 only
Ubuntu 22.04	Python 3.11 only	Python 3.10 only

Crucially, Cloudera Manager 7.13.1 CHF4 release introduces support for Python 3.11. This expansion allows customers to adopt the latest Python version with minimal cluster disruption, benefiting from easier migration, enhanced security, and improved long-term compatibility for Cloudera deployments.

Important: Cloudera Runtime 7.3.1.400 SP2 does not currently support multi-Python support.

Cloudera recommends you to install Python 3.9 or Python 3.11 before upgrading Cloudera Manager to 7.13.1 CHF4 version to ensure smooth transition with minimal downtime. For information about migrating Python for various operating systems, see Migrating Python versions.

Updated the default value for "Region Mover Threads" property for HBase

The default value of Cloudera Manager HBase Configuration Region Mover Threads is changed to 30. This speeds up the rolling restart functionality for HBase.

Use specified users instead of "hive" for Ozone replication-related commands

Starting from Cloudera Manager 7.11.3 CHF15, Ozone commands executed by Ozone replication policies are run by impersonating the users that you specify in the Run as Username and Run on Peer as Username fields in the **Create Ozone replication policy** wizard. The bucket access for OBS-to-OBS replication depends on the user with the access key specified in the fs.s3a.access.key property.

When the source and target clusters are secure, and Ranger is enabled for Ozone, specific permissions are required for Ozone replication to replicate Ozone data using Ozone replication policies. For information about the permissions, see Preparing clusters to replicate Ozone data.

Added Safety Valve for hadoop-metrics2.properties for Ozone roles

Safety Valve for hadoop-metrics2.properties is now available for Ozone roles to enable tuning metrics collection.

Cloudera Manager 7.13.1 Cumulative hotfix 3 (7.13.1.300)

Improved Ranger Admin Diagnostic collection configuration

A new configuration option, ranger.admin.diag.metrics.collection.type, has been introduced. It allows users to specify the type of metrics data to be collected. This replaces the previous behavior of collecting all metrics types by default. Users can now configure the desired metrics type in Ranger, which is then respected by Cloudera Manager's `Send Diagnostic Data` command.

This enhancement reduces the overall data collection time and minimizes the chances of errors, as users can select alternative metrics types if one is slow to collect data or causes issues.

Cloudera Manager 7.13.1 Cumulative hotfix 2 (7.13.1.200)

Floodix Daemon support for Cloudera Manager 7.13.1 CHF2

From Cloudera Manager 7.13.1 CHF2 release, the Floodix Daemon (which uses the client library called Anacrolix/torrent written in Golang) replaces the old Flood Daemon (which uses the Libtorrent protocol and libraries) on all the hosts managed by Cloudera Manager. The purpose of the Floodix daemon and the previous Flood daemon is to efficiently distribute Cloudera Runtime parcels.

With the new Floodix Daemon running, now the Cloudera Manager Server acts as a seeder for all the parcels that are managed across different clusters. In any case, if you want to disable the Cloudera Manager Server acting as a seeder, then perform the steps from How to disable the Cloudera Manager Server acting as a seeder for all the managed parcels.

Rocky Linux 9.4 support for Cloudera Manager

Starting with the Cloudera Manager 7.13.1 CHF2 release, Cloudera Manager provides support for Rocky Linux. This update ensures seamless compatibility with Rocky Linux version 9.4, offering greater flexibility and platform options.

Rocky Linux 9.4 supports only Python 3.9 version in Cloudera Manager 7.13.1 CHF2 release.

Apply java options for each Ozone role separately

Earlier, java options added in the Ozone configurations using the Ozone Java Options configuration are applied to all the Ozone roles together. You cannot add the java options individually to each Ozone role separately.

Now, new configurations such as Ozone Manager Java Options, Ozone SCM Java Options, and so on allows you to add java options to each Ozone role separately. Also, common java options added to the Ozone Java Options configuration apply to all the Ozone roles.

Hive Metastore connection pool metrics in Cloudera Runtime Charts

Starting with Cloudera Runtime 7.3.1.200 SP1 new Hive Metastore charts in Cloudera Manager display connection pool metrics. These charts provide visibility into the connection pools(such as the objectstore, txnhandler) by following pool metrics:

- objectstore.pool.ActiveConnections
- objectstore.pool.IdleConnections
- objectstore.pool.PendingConnections
- objectstore.pool.TotalConnections

This enhancement helps you monitor and optimize Hive Metastore performance.

New KMS Tomcat metrics are available in Cloudera Manager CHF2

New Tomcat container metrics are added in Cloudera Manager CHF2. They can be used by the end users to monitor Tomcat operations. These metrics are available in chart builder and can be used to create a new chart as per the requirement. List of new metrics added are as follows:

- ranger_kms_max_connections
- ranger_kms_active_connections
- ranger_kms_accept_connections
- ranger_kms_connection_timeout
- ranger_kms_connection_keepalive_timeout
- ranger_kms_max_worker_thread_count
- ranger_kms_min_worker_thread_count
- ranger_kms_active_worker_thread_count
- ranger_kms_total_worker_thread_count

Cloudera Manager 7.13.1 Cumulative hotfix 1 (7.13.1.100)

Secure Approach for Passing a Token in Cloudera Manager

You can now securely manage the secret token for the LLM hosting service through Cloudera Manager. Previously, the secret token had to be stored as plain text in Hue’s safety valve configuration. This enhancement improves security and compliance.

For more information, see Secure Approach for Passing a Token in Cloudera Manager.

Replace the Rolling Restart with Restart during ECS upgrade

Enabled the Restart back in ECS, so that we can do a Restart on ECS cluster, services and roles. This will be a combination of Stop and Start operation. Also, the Rolling Restart after the ECS upgrade will be a simple Restart.

Added Services Health Check to the ECS Pre-Upgrade UI

A list of pre-upgrade checks are added that runs after the upgrade version has been chosen. This checklist verifies if your cluster is ready for upgrade.

Support for FIPS cluster in Replication Manager

You can use the FIPS source and target clusters for supported replication policies in Cloudera Base on premises 7.3.1 CHF1 and higher using Cloudera Manager 7.13.1 CHF1 and higher versions. Before you use the FIPS clusters in Replication Manager, ensure that you view the *Support for FIPS clusters* section in Support matrix.

Cloudera Manager 7.13.1

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 7.13.1.500 CHF5

Important: The complete support for Cgroup v2 will start from Cloudera Manager 7.13.1.500 CHF 5 release.

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:

- Cloudera Manager currently does not support Cgroup v2 on **Ubuntu 22.04**. Additionally, Cloudera Manager does not support hybrid Cgroup configurations where both Cgroup v1 and v2 coexist.
- Cloudera does not support Cgroup v2 on RHEL 7, as it is a deprecated operating system. Therefore, no support or testing is provided for cgroups v2 on the RHEL 7 platform.
- For the users using RHEL 9.x with Cloudera Manager version lower than 7.13.1.500 CHF5, must disable Cgroup v2 if already enabled before upgrading to Cloudera Manager 7.13.1.500 CHF5 version as cgroup v2 is not supported with Cloudera Manager version lower than 7.13.1.500 CHF5.
- 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 Cloudera 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 Cloudera Replication Manager

The following changes are available for Iceberg replication policies in Cloudera 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 and its cumulative hotfixes.

Cloudera Manager 7.13.1 CHF5 (7.13.1.500)

There are no new Platform Support Enhancements in this release.

Cloudera Manager 7.13.1 CHF4 (7.13.1.400)

Platform Support Enhancements

RHEL 8.10 FIPS (extended RHEL 8.10 support for FIPS customers)

Cloudera Manager 7.13.1 CHF3 (7.13.1.300)

Platform Support Enhancements

New OS support: RHEL 9.5

Cloudera Manager 7.13.1 CHF2 (7.13.1.200)

Platform Support Enhancements

New OS support: Rocky Linux 9.4

Cloudera Manager 7.13.1 CHF1 (7.13.1.100)

There are no new Platform Support Enhancements in this release.

Cloudera Manager 7.13.1

Platform Support Enhancements

- **New OS support**: RHEL 9.4
- **New Database support**: None
- **New JDK Version**: None

Release Matrix

You must be familiar with the versions of all the cumulative hotfixes for Cloudera Manager 7.13.1.

The Release build number has two parts, the [*Cloudera Manager release version number*] and the [*Cloudera Manager Build number*]. For example, in 7.13.1.100-69, 7.13.1.100 is the release version number and 69 is the build number.

Table 1: Cloudera Manager versions certified with Cloudera Base on premises

Cumulative Hotfix	Cloudera Manager Release Build Number	GBN	Release Date	Download URL
7.13.1 CHF5	7.13.1.500-79	71959111	November 5, 2025	https://archive.cloudera.com/p/cm7/7.13.1.500/
7.13.1 CHF4	7.13.1.400-45	68000784	June 27, 2025	https://archive.cloudera.com/p/cm7/7.13.1.400/
7.13.1 CHF3	7.13.1.300-38	66873372	June 2, 2025	https://archive.cloudera.com/p/cm7/7.13.1.300/
7.13.1 CHF2	7.13.1.200-43	65138596	April 24, 2025	https://archive.cloudera.com/p/cm7/7.13.1.200/
7.13.1 CHF1	7.13.1.100-69	63338448	March 18, 2025	https://archive.cloudera.com/p/cm7/7.13.1.100/
7.13.1	7.13.1.0-333	60343691	December 10, 2025	https://archive.cloudera.com/p/cm7/7.13.1.0/

Table 2: Cloudera Manager versions certified with Cloudera Data Services on premises

Cumulative Hotfix	Cloudera Manager Release Build Number	GBN	Release Date	Download URL
7.13.1 CHF5	7.13.1.501-2	72546494	November 7, 2025	https://archive.cloudera.com/p/cm7/patch/7.13.1.501/
7.13.1 CHF4	7.13.1.402-1	68803632	July 16, 2025	https://archive.cloudera.com/p/cm7/patch/7.13.1.402/
7.13.1 CHF3	7.13.1.303-1	68240225	July 1, 2025	https://archive.cloudera.com/p/cm7/7.13.1.303/

Fixed Issues

Review the list of Cloudera Manager issues that are resolved in Cloudera Manager 7.13.1 and its cumulative hotfixes.

Cloudera Manager 7.13.1 CHF5 (7.13.1.500)

DMX-3659, DMX-3670: Updated row is replicated as inserted row

Previously, during an Iceberg replication policy run, an updated row on the source cluster was replicated as an inserted row. As a result, the number of rows for the table in the source and target clusters did not match after the replication job completed. This issue is now resolved.

DMX-3637: Incorrect report of Iceberg replication

Previously, the number of files transformed during the Iceberg replication policy run was incorrect when multiple tables existed in the Iceberg replication policy. This issue is now resolved

DMX-3627: Iceberg replication policies failed with exception

Previously, some Iceberg replication policies failed with the java.lang.ClassNotFoundException: com.cloudera.repl.iceberg.PathUtil exception. This issue is now resolved.

CDPD-91410: Upgraded external_versions jetty to 9.4.58 for Replication Manager

Upgraded the external_versions jetty to 9.4.58 to fix CVE-2025-5115 for Replication Manager.

OPSAPS-74212: Atlas rolling upgrade related to Zero Downtime Upgrade (ZDU) fails from from 7.1.7 SP3 to 7.3.1.500

The issue causing ZDU failures during upgrades from Cloudera Runtime 7.1.7 SP3 to 7.3.1.500 has been resolved. Previously, the Atlas rolling upgrade was failing because the RoleState for Atlas was not checked, and the upgradeCommand was not set correctly. The updated logic now includes a check for the RoleState of Atlas. Based on this state, the appropriate upgradeCommand is determined and set, ensuring a smooth and reliable rolling upgrade path in ZDU-enabled scenarios.

OPSAPS-74091: Migrate Navigator Encrypt keys to Ranger KMS from KTS configured with HSM

Exporting Navigator Encrypt keys from KTS to Ranger KMS is already available. But if HSM is configured with KTS, this does not work as key's content does not contain the actual key material; it needs to be fetched from HSM first.

Condition has been added to check for HSM setup and accordingly publish a warning log stating Navigator Encrypt keys with HSM cannot be migrated, along with the document link for the steps to migrate.

OPSAPS-74341: NodeManagers might fail to start during the cluster restart after the Cloudera Manager 7.13.1.x upgrade

Cgroup v2 support is enabled in CDP 7.1.9 SP1 CHF5 and higher versions. However, if the user upgrades from Cloudera Manager 7.11.3.x to Cloudera Manager 7.13.1.x, and the environment is using cgroup v2, the NodeManagers might fail to start during the cluster restart after the Cloudera Manager 7.13.1.x upgrade.

This issue is fixed now.

OPSAPS-73038: False-positive port conflict error message displayed in Cloudera Manager

This issue is fixed now. A new health port has been added as a configuration to the Knox configuration. The health topology port can be set with topology port mapping. By setting the new configuration, the checkDeployment script will use the new health port.

OPSAPS-73498: Backport Cloudera Manager side Ranger-Trino integration changes

Trino plugin support in Ranger has been added.

OPSAPS-74276: RockDB JNI library is loaded from the same place to multiple Ozone components

By default, Ozone roles define a separate directory to load RocksDB shared library, and clean up separately from each other on the same host, unless the environment already defines the ROCKSDB_SHAREDLIB_DIR variable via a Safety valve as suggested in the workaround for OPSAPS-67650. After this change, that workaround becomes obsolete. The new directory used reside within directories used by the Cloudera Manager agent to manage the Ozone related processes.

OPSAPS-73628: Impala query profile export to Telemetry Publisher failed due to a 5MB string length limit introduced in Jackson 2.15.0.

The Jackson string length limit was increased to allow exporting large Impala query profiles. Specifically, maxStringLength was set to Integer.MAX_VALUE using StreamReadConstraints, resolving the export failure.

OPSAPS-73709: When Ranger is enabled, Telemetry Publisher fails to export Hive payloads from Data Hub due to missing Ranger client dependencies in the Telemetry Publisher classpath.

This issue has been resolved by adding the necessary dependencies to the classpath.

OPSAPS-73912: In the Cloudera Manager versions 7.13.1, the PROXY settings for the Telemetry Publisher (TP) are not functioning as expected. This may impact the Telemetry Publisher's ability to communicate through a configured proxy.

This issue is resolved by updating the **cdp-sdk-java** JAR to a version that supports proxy settings.

OPSAPS-74392: When creating a compressed archive of an input directory, an open input stream was not closed before a file was deleted. This could lead to filesystem errors, such as the creation of .nfs files.

The issue is now resolved by ensuring the input stream for each file is closed when adding it to the archive.

OPSAPS-73188: Stable Hive queries in a JDK17 environment

Hive queries failed in JDK17 environments because of an InaccessibleObjectException. This was caused by outdated Java options that were not compatible with the new JDK version.

This issue is now resolved by an upgrade handler for Tez that automatically adds the required JDK17 configuration flags. This change automates the manual workaround, ensuring that Hive queries run successfully in JDK17 environments

OPSAPS-73359: Snappy native library loading failure

Snappy native library loading fail in certain cluster configurations. This occurs because Snappy attempts to locate its .so files in /var/lib/hive.

This issue is now fixed.

OPSAPS-72446, OPSAPS-71565, OPSAPS-71566, OPSAPS-73405,OPSAPS-72860: Replication policy runs when the source or target cluster becomes available after it recovers from temporary node failures

Hive replication policies and HBase replication policies can now recover from a temporary node failure on the source or target clusters to continue the replication policy job run. Alternatively, you can also rerun the failed or aborted policies manually.

To ensure that the RemoteCmdWork daemon continues to poll even in case of network failures or if the Cloudera Manager goes down, you can set the remote_cmd_network_failure_max_poll_count=*[*** ENTER REMOTE EXECUTOR MAX POLL COUNT***]* parameter on the target Cloudera Manager Administration Settings page. The actual timeout is provided by a piecewise constant function that is a step function with the following breakpoints: 1 through 11 is 5 seconds, 12 through 17 is 1 minute, 18 through 35 is 2 minutes, 36 through 53 is 5 minutes, 54 through 74 is 8 minutes, 75 through 104 is 15 minutes, and so on. Therefore when you enter 1, the polling continues for 5 seconds after the Cloudera Manager goes down or after a network failure. Similarly when you set it to 75, the polling continues for 15 minutes.

To ensure Replication Manager attempts to recover the RemoteCmdWork daemon on the target cluster, ensure that you set the retry value in the target Cloudera Manager Administration Settings remote_cmd_max_recovery_count parameter, or set it to 0 to turn off the feature. By default, Replication Manager attempts to recover the command twice after the target cluster goes down temporarily.

This issue is now fixed.

OPSAPS-71459: Commands continue to run after Cloudera Manager restart

Previously, some remote replication commands continued to run endlessly even after a Cloudera Manager restart operation. This issue is now fixed.

OPSAPS-72439, OPSAPS-74278, OPSAPS-74265: HDFS and Hive external tables replication policies failed when using custom krb5.conf files

The issue appeared because the custom krb5.conf file was not propagated to the required files. This issue can now be fixed by using a custom Kerberos path before running the replication policies as described in step 13 in Using a custom Kerberos configuration path.

OPSAPS-73645, OPSAPS-73847: Ozone bucket browser does not show the volume buckets

Previously, when you clicked on Next Page on the Cloudera Manager Clusters Ozone Bucket Browser page, and then on a volume name, the volume buckets did not appear if the number of volumes exceeded 26. This issue is now fixed.

OPSAPS-73780: Existing Iceberg replication policies fail after Cloudera Manager is upgraded

When you upgrade the source and target clusters to Cloudera Manager 7.13.1.x but continue to use CDP Private Cloud Base 7.1.9.x version, existing Iceberg replication policies fail because of

compatibility issues. After you upgrade to Cloudera Manager 7.13.1.x, you must also upgrade to Cloudera on premises 7.3.1.x to avoid compatibility issues.

Always ensure that the Cloudera on premises and Cloudera Manager versions are compatible before you run replication policies.

OPSAPS-73906, OPSAPS-73737, OPSAPS-73655, OPSAPS-74061: Cloud replication no longer fails after the delegation token is issued

Previously, the replication policies were failing during incremental replication job runs if you chose the Advanced Setting Delete Policy Delete permanently option during the replication policy creation process.

You can now configure com.cloudera.enterprise.distcp.skip-delegation-token-on-cloud-replication to false in the Cloudera Manager Clusters *HDFS SERVICE* Configuration HDFS Replication Advanced Configuration Snippet (Safety Valve) for core-site.xml advanced configuration snippet to ensure that the HDFS and Hive external table replication policies replicating from an on-premises cluster to cloud do not fail.

When the advanced configuration snippet is set to false, the MapReduce client process obtains the delegation tokens explicitly before it submits the MapReduce job for the replication policy. By default, the advanced configuration snippet is set to true.

OPSAPS-74026: Error appears when you disable Iceberg replication policies

Previously, the Invalid Peer Exception error appeared when you disabled or enabled an Iceberg replication policy on the Cloudera Manager Replication Replication Policies page. This issue is now resolved.

OPSAPS-74040, OPSAPS-74058: Ozone OBS replication fails due to pre-filelisting check failure

During OBS-to-OBS Ozone replication, if the source bucket is a linked bucket, the replication failed during the `Run Pre-Filelisting Check` step, and the error message Source bucket is a linked bucket, however the bucket it points to is also a link appeared, even when the source bucket directly links to a regular, non-linked bucket. The issue is now fixed.

The Ozone OBS-to-OBS replication no longer fails when the source or the target bucket is a linked bucket because the linked bucket resides in the s3v volume and refers to another bucket in s3v or any other volume.

OPSAPS-73158, OPSAPS-73903, OPSAPS-74206: HDFS replication policies fail when the policies pre-fetch the expired Kerberos ticket from the 'sourceTicketCache' file

Previously, Replication Manager pre-fetched the Kerberos ticket from the sourceTicketCache file for the replication policies. Issues appeared when the file contained an expired Kerberos ticket.

To ensure that the replication policies do not pre-fetch the expired Kerberos ticket from the sourceTicketCache file before the replication policy run, Replication Manager now determines whether the sourceTicketCache is current and valid. If its not valid, it continues to fetch tickets until a valid ticket is identified.

If you do not want Replication Manager to perform this step and always fetch a new ticket, you can add USE_SOURCE_PREFETCHED_KERBEROS_PRINCIPAL = false in the Cloudera Manager Clusters HDFS service Configuration HDFS Replication Environment Advanced Configuration Snippet (Safety Valve) advanced configuration snippet.

OPSAPS-73602, OPSAPS-74353: HDFS replication policies to cloud failed with HTTP 400 error

Previously, the HDFS replication policies to cloud were failing after you edited the replication policies in the Cloudera Manager Replication Manager UI . This issue is now fixed.

OPSAPS-74950: Ozone replication policies failed for Cloudera Private Cloud Base 7.1.9 SP1 CHF11 clusters using Cloudera Manager 7.13.1.400

Previously, Ozone replication policies for Ozone linked buckets failed when the Cloudera Private Cloud Base 7.1.9 SP1 CHF11 source or target clusters used Cloudera Manager 7.13.1.400.

To mitigate this issue, use Cloudera Private Cloud Base 7.1.9 SP1 CHF11 source or target clusters with Cloudera Manager 7.13.1.500.

OPSAPS-73572: Add HBase client TLS parameters to the HBase server and Gateway roles

New client-side TLS configuration parameters are added to Cloudera Manager for both the HBase server and Gateway roles. You can now control the TLS client-side configuration settings directly in Cloudera Manager.

Cloudera Manager 7.13.1 CHF4 (7.13.1.400)

OPSAPS-73370: Enhance code to merge compressed Spark event log files

Fixes an issue with unreported metrics in Cloudera Observability when the spark.eventLog.compress property was set to true.

The spark.eventLog.compress property is set to false by default, but enabling it will no longer encoded event logs to fail when processed in Cloudera Observability.

OPSAPS-60642: Host header injection issue on /j_spring_security_check internal endpoint

/j_spring_security_check is internal endpoint which is vulnerable to Host header injection. This issue occurs if the user disabled PREVENT_HOST_HEADER_INJECTION feature flag.

Host header injection: In an incoming HTTP request, web servers often dispatch the request to the target virtual host based on the value supplied in the Host header. Without proper validation of the header value, the attacker can supply invalid input to cause the web server to:

- Dispatch requests to the first virtual host on the list
- Redirect to an attacker-controlled domain
- Perform web cache poisoning
- Manipulate password reset functionality

This issue is resolved now by adding Feature Flag PREVENT_HOST_HEADER_INJECTION to prevent host header injection vulnerability on /j_spring_security_check internal endpoint. This feature flag is by default enabled and it enables additional logic to block potential Host Header Injection attacks targeting the /j_spring_security_check endpoint in Cloudera Manager.

OPSAPS-74019/OPSAPS-72739: Query execution stability with temporary directories

Queries previously failed with an execution error when using a compression library. Although /tmp is a default temporary folder, its use for script execution was blocked due to security restrictions, causing queries to fail.

This issue was resolved by configuring Hive to use a different default temporary folder, /var/lib/hive, instead of /tmp.

OPSAPS-74141: Hive service setup on reused databases

During 7.3.1 base cluster installations, the Hive service setup failed when attempting to validate the Hive Metastore Schema. This happened specifically when the new cluster used a database that had been previously used by an older installation, causing the schema validation to fail due to a version mismatch with the newer Hive components.

This issue was resolved by modifying the Hive script in Cloudera Manager to use the -initOrUpgradeSchema argument instead of -initSchema for the hive_metastore_create_tables command. This change allows the Hive Metastore schema to be properly initialized or upgraded even when connecting to a database that was previously used by an older installation.

OPSAPS-73011: Wrong parameter in the /etc/default/cloudera-scm-server file

In case the Cloudera Manager needs to be installed in High Availability (2 nodes or more as explained here), the parameter CMF_SERVER_ARGS in the /etc/default/cloudera-scm-server file is missing the word "export" before it (on the file there is only CMF_SERVER_ARGS= and not export CMF_SERVER_ARGS=), so the parameter cannot be utilized correctly.

This issue is fixed now.

OPSAPS-72756:The runOzoneCommand API endpoint fails during the Ozone replication policy run

The `/clusters/{clusterName}/runOzoneCommand` Cloudera Manager API endpoint fails when the API is called with the `getOzoneBucketInfo` command. In this scenario, the Ozone replication policy runs also fail if the following conditions are true:

- The source Cloudera Manager version is 7.11.3 CHF11 or 7.11.3 CHF12.
- The target Cloudera Manager is version 7.11.3 through 7.11.3 CHF10 or 7.13.0.0 or later where the feature flag API_OZONE_REPLICATION_USING_PROXY_USER is disabled.

This issue is fixed now.

OPSAPS-72710: Marking the snapshots created by incremental replication policies differently

In the Ozone bucket browser, the snapshots created by an Ozone replication are marked. When the snapshots are deleted, a confirmation modal window appears before the deletion. The restore bucket modal window now displays information about how the restore operation is implemented in Ozone and how this operation affects Ozone replications.

OPSAPS-72447, CDPD-76705: Ozone incremental replication fails to copy renamed directory

Ozone incremental replication using Ozone replication policies succeed but might fail to sync nested renames for FSO buckets.

When a directory and its contents are renamed between the replication runs, the outer level rename synced but did not sync the contents with the previous name.

This issue is fixed now.

OPSAPS-71046: The jstack logs collected on Cloudera Manager 7.11.3 are not in the right format

On viewing the jstack logs in the user cluster, the jstack logs for ozone and other services on Cloudera Manager 7.11.3 and CDP Private Cloud Base 7.1.9 are not in the right format. This issue is fixed now.

OPSAPS-65377: Cloudera Manager - Host Inspector not finding Psycopg2 on Ubuntu 20 or Redhat 8.x when Psycopg2 version 2.9.3 is installed.

Host Inspector fails with Psycopg2 version error while upgrading to Cloudera Manager 7.13.1.x versions. When you run the Host Inspector, you get an error Not finding Psycopg2, even though it is installed on all hosts. This issue is fixed now.

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

The solrconfig.xml and schema.xml files are updated for the existing atlas_configs directory to ensure backward compatibility for existing customers. This included downloading the current Atlas configuration from Solr, applying the necessary updates, and then overwriting the modified configuration back to Solr. This issue is fixed now and Atlas uses the correct configuration directory in /var/run/cloudera-scm-agent/process/151-atlas-ATLAS_SERVER/solrconf.xml. New clusters already use the updated configuration directory.

OPSAPS-74147: Atlas rolling upgrade related to Zero Downtime Upgrade (ZDU) fails from from 7.1.7 SP3 to 7.3.1.400

The issue causing ZDU failures during upgrades from Cloudera Runtime 7.1.7 SP3 to 7.3.1.400 has been resolved. Previously, the Atlas rolling upgrade was failing because the RoleState for Atlas was not checked, and the upgradeCommand was not set correctly.

OPSAPS-73174: Autoscaling fails when any of the RM hosts are down

When a master node hosting RM abruptly goes down, Cloudera Manager can proceed with the NM commission/decommission command-flow.

OPSAPS-73780: Existing Iceberg replication policies fail after Cloudera Manager is upgraded

Existing Iceberg replication policies fail when the source and target clusters are upgraded to Cloudera Manager 7.13.1.x which use CDP Private Cloud Base 7.1.9.x version. The replication

policies fail because of compatibility issues. You must ensure that the CDP versions and Cloudera Manager versions are compatible before you run replication policies. For example, use Cloudera on premises 7.3.1.x with Cloudera Manager 7.13.1.x. This issue is resolved.

Cloudera Manager 7.13.1 CHF3 (7.13.1.300)

OPSAPS-73225: Cloudera Manager Agent reporting inactive/failed processes in Heartbeat request

As part of introducing Cloudera Manager 7.13.x, some changes were done to the Cloudera Manager logging, eventually causing Cloudera Manager Agent to report on inactive/stale processes during Heartbeat request.

As a result, the Cloudera Manager servers logs are getting filled rapidly with these notifications though they do not have impact on service.

In addition, with adding the support for the Observatory feature, some additional messages were added to the logging of the server. However, in case the customer did not purchase the Observatory feature, or the telemetry monitoring is not being used, these messages (which appears as "TELEMETRY_ALTUS_ACCOUNT is not configured for Otelcol" are filling the server logs and preventing proper follow-up on the server activities).

This issue is fixed now.

OPSAPS-72270: Start ECS command fails on uncordon nodes step

Ensure the `kube-apiserver` is up and running for at least 60 seconds before proceeding with the uncordon step, and use the correct target node name, not just the name of the node where the uncordon command is executed.

This issue is fixed now.

OPSAPS-72978: The getUsersFromRanger API parameter truncates the user list after 200 items

The Cloudera Manager API endpoint v58/clusters/*[***CLUSTER***]*/services/*[***SERVICE***]*/commands/getUsersFromRanger API endpoint no longer truncates the list of returned users at 200 items.

Cloudera Manager 7.13.1 CHF2 (7.13.1.200)

OPSAPS-72809: Ranger policy script for Knox fails due to double quotation marks

The Ranger policy script for Knox (setupRanger.sh) fails, because the CSD_JAVA_OPTS parameters are enclosed by double quotation marks in the script. The issue is fixed now.

OPSAPS-72795: Do not allow multiple Ozone services in a cluster

It is possible to configure multiple Ozone services in a single cluster which can cause irreversible damage to a running cluster. So, this fix allows you to install only one Ozone service in a cluster.

OPSAPS-72767: Install Oozie ShareLib Cloudera Manager command fails on FIPS and FedRAMP clusters

The Install Oozie ShareLib command using Cloudera Manager fails to execute on FIPS and FedRAMP clusters. This issue is fixed now.

OPSAPS-72323: Cloudera Manager UI is down with bootstrap failure due to ConfigGenExecutor throwing exception

This issue is fixed now.

OPSAPS-71566: The polling logic of RemoteCmdWork goes down if the remote Cloudera Manager goes down

When the remote Cloudera Manager goes down or when there are network failures, the RemoteCmdWork stops to poll. To ensure that the daemon continues to poll even when there are network failures or if the Cloudera Manager goes down, you can set the remote_cmd_network_failure_max_poll_count=[*** ENTER REMOTE EXECUTOR MAX POLL

COUNT***] parameter on the Cloudera Manage Administration Settings page. Note that the actual timeout is provided by a piecewise constant function (step function) where the breakpoints are: 1 through 11 is 5 seconds, 12 through 17 is 1 minute, 18 through 35 is 2 minutes, 36 through 53 is 5 minutes, 54 through 74 is 8 minutes, 75 through 104 is 15 minutes, and so on. Therefore when you enter 1, the polling continues for 5 seconds after the Cloudera Manager goes down or after a network failure. Similarly when you set it 75, the polling continues for 15 minutes.

OPSAPS-67197: Ranger RMS server shows as healthy without service being accessible

Being a Web service, Ranger RMS might not be initialized due to other issues causing RMS to be inaccessible. But Ranger RMS service was still shown as healthy, because Cloudera Manager only monitors Process Identification Number (PID).

This issue is fixed now. Added the health status canary support for Ranger RMS service which connects to RMS after some specific intervals and shows alert on the Cloudera Manager UI if RMS is not reachable.

OPSAPS-71933: Telemetry Publisher is unable to publish Spark event logs to Cloudera Observability when multiple History Servers are set up in the Spark service.

This issue is now resolved by adding the support for multiple Spark History Server deployments in Telemetry Publisher.

OPSAPS-71623: Some Spark jobs are missing from the Workload XM interface. In the Telemetry Publisher logs for these Spark jobs, the error message java.lang.IllegalArgumentException: Wrong FS for Datahub cluster is displayed.

The issue is resolved by addressing Telemetry Publisher failures during the processing of Yarn logs.

Cloudera Manager 7.13.1 CHF1 (7.13.1.100)

OPSAPS-71669: The Continue option is disabled on the Static Service Pools Review page, affecting the functionality of Static Service Pools

The minimum and maximum I/O weight values for Cgroup v2 were incorrectly set to 100 and 1000, respectively, in Cloudera Manager 7.13.1.0. According to official Cgroup v2 documentation, the valid range should be 1 to 10,000. Due to this incorrect configuration range, the Continue option on the **Static Service Pools Review** page was disabled, preventing users from proceeding with pool configuration.

This issue might occur on clusters running Cloudera Manager 7.13.1.0 with Cgroup v2 resource management when configuring or reviewing Static Service Pools. After upgrading to Cloudera Manager 7.13.1.100 CHF-1, this issue no longer occurs.

This issue is fixed now. The minimum and maximum I/O weight values for Cgroup v2 have been corrected to 1 and 10,000, respectively, in accordance with the official Cgroup v2 specification. This fix ensures that the Continue option on the **Static Service Pools Review** page is now enabled as expected.

OPSAPS-72369: Update snapshot default configuration for enabling ordered snapshot deletion

This issue is now resolved by changing the default configuration value on Cloudera Manager.

OPSAPS-72215: ECS CM UI Config for docker cert CANNOT accept the new line - unable to update new registry cert in correct format

Currently there is no direct way to update the external docker certificate in the UI for ECS because newlines are removed when the field is saved. Certs can be uploaded by adding '\n' character for newline now. When user wants to update docker cert through Cloudera Manager UI config. User need to add '\n' to specify a newline character in the certificate. Example:

```
-----BEGIN CERTIFICATE-----\
nMIIERTCCAy2gAwIBAgIUIL8o1MjD5he7nZKKa/C8rx9uPjcwDQYJKoZIhvcNAQEL
\n
```

```
BQAwXTELMAkGA1UEBhMCVVMxEzARBgNVBAgMCkNhbGlmb3JuaWExEzARBgNVBAcM
\nClNhbnRhQ2xhcmExETA
PBgNVBAoMCENsb3VkZXJhMREwDwYDVQQLDAhDbG91ZGVy\nYTAeFw0yNDAzMTExMj
U5NDVaFw0zNDAzMDkxMjU5NDVaMF0x
CzAJBgNVBAYTAlVT\nMRMwEQYDVQQIDApDYWxpZm9ybmlhMRMwEQYDVQQHDApTYW5
0YUNsYXJhMREwDwYD\nVQQKDA
hDbG91ZGVyYTERMA8GA1UECwwIQ2xvdWRlcmEwggEiMA0GCSqGSIb3DQEB\nAQ
UAA4IBDwAwggEKAoIBAQDcuxGszWmzVnWCwDICnlxUBtO
+Ps2jxQ7C7kIj\nTHTaQ2kGl/ZzQOJBpYT/jFmiQGPSKb4iLSxed+Xk5xAOkNWDIL
+Hlf5txjkw/FTf\nHiyWep9DaQDF07M/Cl3nb8JmpRyA5fKYpVbJAFIEXOhT
xrcnH/4o5ubLM7mHVXwY\nafoPD5AuiOD/I+xxmqb/x+fKtHzY1eEzDb2vjjDJBR
qxpHvg/S4hHsgZJ7wU7wg+\nPk4uPV3O83h9NI+b4SOwXunuKRCCh4dRKm8/Qw4f
7tDFdCA
IubvO1AGtfyJJp9xR\npMIjhIuna1K2TnPQomdoIy/KqrFFzVaHevyinEnRLG2NA
gMBAAGjgfwwgfkwHQYD\nVR0OBBYEFHWX21/BhL5J5kNpxmb8F
mDchlmBMIGaBgNVHSMEgZIwgY+AFHWX21/B\nhL5J5kNpxmb8FmDchlmBoWGkXzBd
MQswCQYDVQQGEwJVUzETMBEGA1UECAwKQ2Fs\naWZvcm5pYTET
MBEGA1UEBwwKU2FudGFDbGFyYTERMA8GA1UECgwIQ2xvdWRlcmEx\nETAPBgNVBAs
MCENsb3VkZXJhghQgvyjUyMPmF7udkopr8LyvH24+NzAMBgNVHRM
E\nBTADAQH/MAsGA1UdDwQEAwIC/DAPBgNVHREECDAGhwQKgW26MA8GA1UdEgQI
MAaH\nBAqBbbowDQYJKoZIhvcNAQELBQADggEBAMks+sY+ETaPzFLg2PolUT
4GeXEqnGl\nSmZzIkiA6l2DCYQD/7mTLd5Ea63oI78fxatRnG5MLf5aHVLs4W+W
YhoP6B7HLPUo\nNGPJviRBHtUDRYVpD5Q0hhQtHB4Q1H+sgrE53VmbIQqLPOAxvp
M//oJCFDT8N
bOI\n+bTJ48N34ujosjNaiP6x09xbzRzOnYd6VyhZ/pgsiRZ4qlZsVyv1TImP9Vp
HcC7P\nukxNuBdXBS3jEXcvEV1Eq4Di+z6PIWoPIHUunQ9P0akYEvbXu
L88knM5FNhS6YBP\nGd91KkGdz6srRIVRiF+XP0e6IwZC70kkWiwf8vX/CuR64Z
Qxc30ot70=\n-----END CERTIFICATE-----\n
```


OPSAPS-72662: UIDs (User IDs) conflicts for the kubernetes containers as the Kubernetes containers use the user ID - 1001 which is a pretty common UID in a Unix environment.

This issue is fixed now by using a large UID such as 1000001 to reduce UID conflicts.

Using large UIDs (User IDs) for Kubernetes containers is a recommended security practice because it helps minimize the risk of a container compromising the host system. By assigning a high UID, it reduces the chances of conflicts with existing user accounts on the host, particularly if the container is compromised and attempts to access host files or escalate privileges. In essence, a large UID ensures the container operates with restricted permissions on the host system. Therefore, when creating the CLI pod in Cloudera Manager, the runAsUser value should be set to an integer greater than 1,000,000. To avoid UID conflicts, it is advisable to use a UID such as 1000001.

Important: Exception case: Where some pods need a specific UID such as the embedded DB and they do not have this change.

OPSAPS-72559: Incorrect error messages appear for Hive ACID replication policies

Replication Manager now shows correct error messages for every Hive ACID replication policy run on the Cloudera Manager Replication Manager Replication Policies Actions Show History page as expected. This issue is fixed now.

OPSAPS-72509: Hive metadata transfer to GCS fails with ClassNotFoundException

Hive external table replication policies from an on-premises cluster to cloud failed during the `Transfer Metadata Files` step when the target is on Google Cloud and the source Cloudera Manager version is 7.11.3 CHF7, 7.11.3 CHF8, 7.11.3 CHF9, 7.11.3 CHF9.1, 7.11.3 CHF10, or 7.11.3 CHF11. This issue is fixed.

OPSAPS-72559: Incorrect error messages appear for Hive ACID replication policies

Replication Manager now shows correct error messages for every Hive ACID replication policy run on the Cloudera Manager Replication Manager Replication Policies Actions Show History page as expected.

OPSAPS-72558, OPSAPS-72505: Replication Manager chooses incorrect target cluster for Iceberg, Atlas, and Hive ACID replication policies

When a Cloudera Manager instance managed multiple clusters, Replication Manager picked the first cluster in the list as the Destination during the Iceberg, Atlas, and Hive ACID replication policy creation process, and the Destination field was non-editable. You can now edit the replication policy to change the target cluster in these scenarios.

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

Replication Manager now skips the replicated files during subsequent Ozone replication policy runs after you add the following key-value pairs in Cloudera Manager Clusters *OZONE SERVICE* Configuration Ozone Replication Advanced Configuration Snippet (Safety Valve) for core-site.xml :

- com.cloudera.enterprise.distcp.ozone-schedules-with-unsafe-equality-check = *[***ENTER COMMA-SEPARATED LIST OF OZONE REPLICATION POLICIES' ID OR ENTER ALL TO APPLY TO ALL OZONE REPLICATION POLICIES***]*

  The advanced snippet skips the already replicated files when the relative file path, file name, and file size are equal and ignores the modification times.

  

  Caution: Usage of this advanced snippet might lead to data loss. For example, if you modified a file on the source or target cluster and the file size remains the same, the advanced snippet ignores the file during the replication run.
- com.cloudera.enterprise.distcp.require-source-before-target-modtime-in-unsafe-equality-check = *[***ENTER TRUE OR FALSE***]*

When you add both the key-value pairs, the subsequent Ozone replication policy runs skip replicating files when the matching file on the target has the same relative file path, file name, file size and the source file's modification time is less or equal to the target file modification time.

OPSAPS-72214: Cannot create a Ranger replication policy if the source and target cluster names are not the same

You could not create a Ranger replication policy if the source cluster and target cluster names were not the same. This issue is fixed.

OPSAPS-71853: The Replication Policies page does not load the replication policies' history

When the sourceService is null for a Hive ACID replication policy, the Cloudera Manager UI fails to load the existing replication policies' history details and the current state of the replication policies on the **Replication Policies** page. This issue is fixed now.

OPSAPS-72181: Currently Apply Host Template checks for active command on the service, if the active command is taking time (like a long-running replication command) then Apply Host Template operation will also get delayed.

This issue is fixed now for certain scenario like when host template has only gateway role then the Apply Host Template operation will not check for active command on service. If host template has other roles than gateway then the behaviour remains same. Apply Host Template with gateway roles only will not wait for any active service command.

OPSAPS-72249: Oozie database dump fails on JDK17

Oozie database dump and load commands couldn't be executed from Cloudera Manager with JDK 17. This issue is fixed now.

OPSAPS-72276: Cannot edit Ozone replication policy if the MapReduce service is stale

You could not edit an Ozone replication policy in Replication Manager if the MapReduce service did not load completely. This issue is fixed.

OPSAPS-71932: Ranger HDFS plugin resource lookup issue

For JDK 17 Isilon cluster, user was not able to create a new policy under cm_hdfs. The connection was failing with the following error message:

```
cannot access class sun.net.util.IPAddressUtil
```

The issue is fixed now. Added sun.net.util package to Ranger Admin java opts for JDK 17.

OPSAPS-71907: Solr auditing URL changed port

The Solr auditing URL generated for Ranger plugin services in the data hub cluster is correct when both the local ZooKeeper and the data lake ZooKeeper have ssl_enabled enabled. However, if the ssl_enabled parameter is disabled on the local ZooKeeper in data hub, the Solr auditing URL changed the port to use 2181.

The fix fetches the Solr auditing URL from the data context of data lake on data hub, resolving the issue where, if the ZooKeeper ssl_enabled parameter is disabled, Solr auditing uses port 2181; a rare, corner-case occurrence.

OPSAPS-71666: Replication Manager uses the required property values in the "ozone_replication_core_site_safety_valve " in the source Cloudera Manager during Ozone replication policy run

During an Ozone replication policy run, Replication Manager obtains the required properties and its values from the ozone_replication_core_site_safety_valve. It then adds the new properties and its values and overrides the value for existing properties in the core-site.xml file. Replication Manager uses this file during the Ozone replication policy run.

Tip: Ozone service uses the core-site.xml file for its activities.

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

Ranger replication policy failed during the transform step because of incorrect source to destination service name mapping. This issue is fixed now.

OPSAPS-71642: GflagConfigFileGenerator is removing the = sign in the Gflag configuration file when the configuration value passed is empty in the advanced safety valve

If the user adds file_metadata_reload_properties configuration in the advanced safety valve with = sign and empty value, then the GflagConfigFileGenerator is removing the = sign in the Gflag configuration file when the configuration value passed is empty in the advanced safety valve.

This issue is fixed now.

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

When the `ozone_replication_core_site_safety_valve` advanced configuration snippet is set to its default value, Replication Manager does not read its value during the Ozone replication policy run. To mitigate this issue, the default value of `ozone_replication_core_site_safety_valve` has been set to an empty value. If you have set any key-value pairs for `ozone_replication_core_site_safety_valve`, then these values are written to core-site.xml during the Ozone replication policy run.

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

The OBS-to-OBS Ozone replication policy jobs failed when the S3 property values for `fs.s3a.endpoint, fs.s3a.secret.key`, and `fs.s3a.access.key` were empty in `Ozone Service Advanced Configuration Snippet (Safety Valve) for ozone-conf/ozone-site.xml` even when these properties were defined in `Ozone Replication Advanced Configuration Snippet (Safety Valve) for core-site.xml`. This issue is fixed.

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. This issue is fixed now.

OPSAPS-71093: Validation on source for Ranger replication policy fails

The Cloudera Manager page would be logged out automatically when you created a Ranger replication policy. This is because the source cluster did not support the getUsersFromRanger or getPoliciesFromRanger API requests. The issue is fixed now, and the required validation on the source completes successfully as expected.

OPSAPS-70848: Hive external table replication policies succeed when the source cluster uses Dell EMC Isilon storage

During the Hive external table replication policy run, the replication policy failed at the Hive Rep lication Export step. This issue is fixed now.

OPSAPS-70822: Save the Hive external table replication policy on the ‘Edit Hive External Table Replication Policy’ window

Replication Manager saves the changes as expected when you click **Save Policy** after you edit a Hive replication policy. To edit a replication policy, you click Actions Edit Configuration for the replication policy on the **Replication Policies** page.

OPSAPS-70721: QueueManagementDynamicEditPolicy is not enabled with Auto Queue Deletion enabled

Whenever Auto Queue Deletion is enabled, the QueueManagementDynamicEdit policy is not enabled. This issue is fixed now and when there are no applications running in a queue, then its capacity is set to zero.

OPSAPS-70449: After creating a new Dashboard from the Cloudera Manager UI, the Chart Title field was allowing Javascript as input

In Cloudera Manager UI, while creating a new plot object, a Chart Title field allows Javascript as input. This allows the user to execute a script, which results in an XSS attack. This issue is fixed now.

OPSAPS-69782: Exception appears if the peer Cloudera Manager's API version is higher than the local cluster's API version

HBase replication using HBase replication policies in CDP Public Cloud Replication Manager between two Data Hubs/COD clusters succeed as expected when all the following conditions are true:

- The destination Data Hub/COD cluster’s Cloudera Manager version is 7.9.0-h7 through 7.9.0-h9 or 7.11.0-h2 through 7.11.0-h4, or 7.12.0.0.
- The source Data Hub/COD cluster's Cloudera Manager major version is higher than the destination cluster's Cloudera Manager major version.
- The Initial Snapshot option is chosen during the HBase replication policy creation process and/ or the source cluster is already participating in another HBase replication setup as a source or destination with a third cluster.

OPSAPS-69622: Cannot view the correct number of files copied for Ozone replication policies

The last run of an Ozone replication policy does not show the correct number of the files copied during the policy run when you load the Cloudera Manager Replication Manager Replication Policies page after the Ozone replication policy run completes successfully. This issue is fixed now.

OPSAPS-72143: Atlas replication policies fail if the source and target clusters support FIPS

The Atlas replication policies fail during the `Exporting atlas entities from remote atlas service` step if the source and target clusters support FIPS. This issue is fixed now.

OPSAPS-67498: The Replication Policies page takes a long time to load

To ensure that the Cloudera Manager Replication Manager Replication Policies page loads faster, new query parameters have been added to the internal policies that fetch the REST APIs for the page which improves pagination. Replication Manager also caches internal API responses to speed up the page load.

OPSAPS-65371: Kudu user was not part of the cm_solr RANGER_AUDITS_COLLECTION policy

Kudu user was not part of the default policy of cm_solr, which prevented to write any Kudu audit logs on Ranger Admin untill Kudu user was manually added to the policy.

The issue is fixed now. Added Kudu user to default policy for cm_solr - RANGER_AUDIT S_COLLECTION, so that Kudu user does not need to be added manually to write audits to Ranger Admin.

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_pyspar k_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 Runtime 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_RO OT_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-72153: Invalid signature when trying to create tags in Atlas through Knox

Atlas, SMM UI, and SCHEMA-REGISTRY throw 500 error in FIPS environment.

This issue is fixed now.

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 Cloudera on 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 Cloudera Base on premises cluster using the Isilon service and a target cloud storage bucket for an HDFS replication policy in Cloudera Base on premises 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-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.safelog
ic.cryptocomply.fips.core ' \
'--add-exports=java.base/sun.security.provider=bctls --module-
path=/cdep/extra_jars ' \
'-Dcom.safelogic.cryptocomply.fips.approved_only=true -Djdk.tl
s.ephemeralDHKeySize=2048 ' \
'-Dorg.bouncycastle.jsse.client.assumeOriginalHostName=true -D
jdk.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 Cloudera Base on premises 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 Cloudera Base on premises 7.3.1 release. Therefore added validation and confirmation messages to the Cloudera Manager upgrade wizard to alert the user to migrate KEYTRUSTEE_SERVER keys to RANGER_KMS before upgrading to Cloudera Base on premises 7.3.1 release.

OPSAPS-70656: Remove KEYTRUSTEE_SERVER & RANGER_KMS_KTS from Cloudera Manager for UCL

The Keytrustee components - KEYTRUSTEE_SERVER and RANGER_KMS_KTS services are not supported starting from the Cloudera Base on premises 7.3.1 release. These services cannot be installed or managed with Cloudera Manager 7.13.1 using Cloudera Base on premises 7.3.1.

OPSAPS-67480: In CDP 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.

DMX-3364: Drop table operation works incorrectly during Iceberg replication

A replicated table was dropped automatically in the target cluster during a subsequent policy run after you dropped the table in the source cluster, and then edited the replication policy to remove the table and added another table to the replication policy. This issue is resolved.

Known Issues

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

Known issues identified in Cloudera Manager 7.13.1 CHF5 (7.13.1.500)

DMX-4540: Incomplete logs for transfer step during Iceberg replication policy job run

7.13.1.500

During the transfer step of the Iceberg replication policy job run, the UI displays a message, including the progress information, encountered errors, and other relevant details. However, this message content is not captured in the results.json file. The file only captures the transfer step

start and end timestamps, and whether the step completed successfully. This discrepancy makes troubleshooting challenging for errors occurred during the transfer step.

None

OPSAPS-74066, OPSAPS-74547: DataHub high memory consumption on Hiveserver load for JDK 17

7.13.1.500

In upgraded DataHub deployments, HiveServer might fail to start due to memory overallocation. This occurs because Cloudera Manager does not account for memory already assigned to Management Service roles when allocating memory for cluster roles. This issue is fixed in fresh installations of Cloudera Manager 7.13.1.500. The updated algorithm now correctly reallocates memory across all roles during cluster setup.

To resolve this issue, use the following API to manually trigger the Cloudera Manager memory allocation algorithm on the host where both HiveServer and management roles are running, and restart cluster to apply updated memory configs.:

```
API Endpoint: POST /api/v57/hosts/reallocateMemory
```

Include the host name (the host where HiveServer and management roles run) in the API request body. This ensures that memory assignments are recomputed correctly, taking all roles on the host into account.

Important:

Cloudera Manager always assigns at least the minimum requested memory for every role, even if this results in allocating more memory than the host's available physical memory.

After running the above API, if it still has overallocation and can't reduce memory allocation further for any role on the node, it will still lead to failures. For example, if the Roles are not starting due to high memory requests, then manually reduce the heap size configuration for the affected role.

Known issues identified in Cloudera Manager 7.13.1 CHF4 (7.13.1.400)

OPSAPS-74668: ozone.snapshot.deep.cleaning.enabled and ozone.snapshot.ordered.deletion.enabled configs are missing with Cloudera 7.1.9 SP1 CHF and Cloudera Manager 7.13.1

7.13.1.400 and 7.13.1.500

Two Ozone Manager configs are missing while using Cloudera 7.1.9 SP1 CHF after upgrading Cloudera Manager version from 7.11.3 to 7.13.1.400.

If you are using Cloudera 7.1.9 SP1 CHF, before upgrading Cloudera Manager version from 7.11.3 to 7.13.1.400, add the following configs to Ozone Manager Advanced Configuration Snippet (Safety Valve) for ozone-conf/ozone-site.xml so that Ozone Manager does not miss important config after the Cloudera Manager upgrade:

```
<property>
<name>ozone.snapshot.deep.cleaning.enabled</name>
<value>false</value>
</property>

<property>
<name>ozone.snapshot.ordered.deletion.enabled</name>
<value>true</value>
</property>
```

OPSAPS-73038: False-positive port conflict error message displayed in Cloudera Manager

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

7.13.1.500

Cloudera Manager might display a false-positive error message: Port conflict detected: 8443 (Gateway Health HTTP Port) is also used by: Knox Gateway during cluster installations. The warning does not cause actual installation failures.

None

OPSAPS-74950: Ozone replication policies fail for Cloudera Private Cloud Base 7.1.9 SP1 CHF11 clusters using Cloudera Manager 7.13.1.400

7.13.1.400

7.13.1.500

Ozone replication policies for Ozone linked buckets fail when the Cloudera Private Cloud Base 7.1.9 SP1 CHF11 source or target clusters use Cloudera Manager 7.13.1.400.

Use Cloudera Private Cloud Base 7.1.9 SP1 CHF11 clusters with Cloudera Manager 7.13.1.500.

OPSAPS-72439: HDFS and Hive external tables replication policies fail when using custom “krb5.conf” files for Cloudera Private Cloud Base 7.1.9 SP1 CHF11 clusters using Cloudera Manager 7.13.1.400

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

7.13.1.500

The issue appears when the custom krb5.conf file is not propagated to the required files, and you are using Cloudera Private Cloud Base 7.1.9 SP1 CHF11 source or target clusters with Cloudera Manager 7.13.1.400.

Use Cloudera Private Cloud Base 7.1.9 SP1 CHF11 clusters with Cloudera Manager 7.13.1.500, and complete the instructions in step 13 in Using a custom Kerberos configuration path.

OPSAPS-71459: Commands continue to run after Cloudera Manager restart

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

7.13.1.500

If you are using Cloudera Private Cloud Base 7.1.9 SP1 CHF11 source or target clusters with Cloudera Manager 7.13.1.400, remote replication commands continue to run endlessly even after a Cloudera Manager restart operation.

Use Cloudera Private Cloud Base 7.1.9 SP1 CHF11 clusters with Cloudera Manager 7.13.1.500.

OPSAPS-73158, OPSAPS-74206: HDFS replication policies fail when the policies prefetch the expired Kerberos ticket from the 'sourceTicketCache' file for Cloudera Private Cloud Base 7.1.9 SP1 CHF11 clusters using Cloudera Manager 7.13.1.400

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

7.13.1.500

If you are using Cloudera Private Cloud Base 7.1.9 SP1 CHF11 source or target clusters with Cloudera Manager 7.13.1.400, Replication Manager pre-fetches the Kerberos ticket from the sourceTicketCache file for the replication policies. Issues appear when the file contains an expired Kerberos ticket.

Use Cloudera Private Cloud Base 7.1.9 SP1 CHF11 clusters with Cloudera Manager 7.13.1.500.

OPSAPS-73405, OPSAPS-71565, OPSAPS-72860, OPSAPS-72859: Replication policies fail even after the source or target cluster becomes available after it recovers from temporary node failures for Cloudera Private Cloud Base 7.1.9 SP1 CHF11 clusters using Cloudera Manager 7.13.1.400

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

7.13.1.500

If you are using Cloudera Private Cloud Base 7.1.9 SP1 CHF11 source or target clusters with Cloudera Manager 7.13.1.400, Hive replication policies and HBase replication policies fail even after the source or target cluster recovers from a temporary node failure.

Use Cloudera Private Cloud Base 7.1.9 SP1 CHF11 clusters with Cloudera Manager 7.13.1.500.

OPSAPS-73655, OPSAPS-73737: Cloud replication fails even after the delegation token is issued for Cloudera Private Cloud Base 7.1.9 SP1 CHF11 clusters using Cloudera Manager 7.13.1.400

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

7.13.1.500

If you are using Cloudera Private Cloud Base 7.1.9 SP1 CHF11 source or target clusters with Cloudera Manager 7.13.1.400, the replication policies fail during an incremental replication run if you chose the Advanced Setting Delete Policy Delete permanently option during the replication policy creation process.

Use Cloudera Private Cloud Base 7.1.9 SP1 CHF11 clusters with Cloudera Manager 7.13.1.500.

OPSAPS-74040, OPSAPS-74058: Ozone OBS replication fails due to pre-filelisting check failure for Cloudera Private Cloud Base 7.1.9 SP1 CHF11 clusters using Cloudera Manager 7.13.1.400

7.13.1.400

7.13.1.500

If you are using Cloudera Private Cloud Base 7.1.9 SP1 CHF11 source or target clusters with Cloudera Manager 7.13.1.400 and the source bucket is a linked bucket, then the replication fails during the `Run Pre-Filelisting Check` step for OBS-to-OBS Ozone replication, and the error message Source bucket is a linked bucket, however the bucket it points to is also a link appears. This issue appears even when the source bucket is directly linked to a regular, non-linked bucket.

Use Cloudera Private Cloud Base 7.1.9 SP1 CHF11 clusters with Cloudera Manager 7.13.1.500.

OPSAPS-73602, OPSAPS-74353: HDFS replication policies to cloud fails with HTTP 400 error for Cloudera Private Cloud Base 7.1.9 SP1 CHF11 clusters using Cloudera Manager 7.13.1.400

7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

7.13.1.500

If you are using Cloudera Private Cloud Base 7.1.9 SP1 CHF11 source or target clusters with Cloudera Manager 7.13.1.400, the HDFS replication policies to cloud fail after you edit the replication policies in the Cloudera Manager Replication Manager UI .

Use Cloudera Private Cloud Base 7.1.9 SP1 CHF11 clusters with Cloudera Manager 7.13.1.500.

OPSAPS-73645, OPSAPS-73847: Ozone bucket browser does not show the volume buckets for Cloudera Private Cloud Base 7.1.9 SP1 CHF11 clusters using Cloudera Manager 7.13.1.400

7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

7.13.1.500

If you are using Cloudera Private Cloud Base 7.1.9 SP1 CHF11 source or target clusters with Cloudera Manager 7.13.1.400, the volume buckets do not appear if the number of volumes exceed 26, when you click on Next Page on the Cloudera Manager Clusters *OZONE SERVICE* Bucket Browser page and then on a volume name.

Use Cloudera Private Cloud Base 7.1.9 SP1 CHF11 clusters with Cloudera Manager 7.13.1.500.

Known issues identified in Cloudera Manager 7.13.1 CHF3 (7.13.1.300)

RELENG-27000: Proper link for bigtop-detect-javahome is missing when using CDP Private Cloud Base 7.1.9 SP1 CHF5 with Cloudera Manager 7.13.1 CHF3.

When using Cloudera Manager 7.13.1 CHF3 with CDP Private Cloud Base 7.1.9 SP1 CHF5 results in inappropriate bigtop-detect-javahome link.

Create a link under /opt/cloudera/parcels/CDH/bin/bigtop-detect-javahome that points to /opt/cloudera/parcels/CDH/lib/bigtop-utils/bigtop-detect-javahome. For example:

```
ln -s /opt/cloudera/parcels/CDH/lib/bigtop-utils/bigtop-detect-javahome /opt/cloudera/parcels/CDH/bin/bigtop-detect-javahome
```

Known issues identified in Cloudera Manager 7.13.1 CHF2 (7.13.1.200)

There are no new known issues identified in this release.

Known issues identified in Cloudera Manager 7.13.1 CHF1 (7.13.1.100)

This section lists the known issues that are identified in this release:

CDPD-79725: Hive fails to start after Datahub restart due to high memory usage

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

After restarting the Cloudera Data hub, the services appears to be down in the Cloudera Manager UI. The Cloudera Management Console reports a node failure error for the master node.

The issue is caused by high memory usage due to the G1 garbage collector on Java 17, leading to insufficient memory issues and thereby moving the Cloudera clusters to an error state.

Starting with Cloudera 7.3.1.0, Java 17 is the default runtime instead of Java 8, and its memory management increases memory usage, potentially affecting system performance. Clusters might report error states, and logs might show insufficient memory exceptions.

To mitigate this issue and prevent startup failures after a Datahub restart, you can perform either of the following actions, or both:

- Reduce the Java heap size for affected services to prevent nodes from exceeding the available memory.
- Increase physical memory for on cloud or on-premises instances running the affected services.

OPSAPS-74370: Knox's Save Alias - IDBroker command fails due to missing variable declaration

7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

Users trying to create IDBroker aliases through the Cloudera Manager UI face issues in Cloudera Manager 7.13.1 using CDP 7.1.9.

The alias(es) can be created using the Knox CLI:

1. ssh to Knox host.
2. export KNOX_GATEWAY_DATA_DIR="/var/lib/knox/idbroker/data"; export KNOX_GATEWAY_CONF_DIR="/var/lib/knox/idbroker/conf"
3. /opt/cloudera/parcels/CDH/lib/knox/bin/knoxcli.sh create-alias <ALIAS_NAME> --cluster <CLUSTER_NAME> --value <ALIAS_VALUE>
4. Verify the addition using /opt/cloudera/parcels/CDH/lib/knox/bin/knoxcli.sh list-alias --cluster <CLUSTER_NAME>

For HA deployments, users must do it on every Knox hosts (whereas the Save Alias command applies the change to all hosts automatically).

Known issues in Cloudera Manager 7.13.1

OPSAPS-70403: Custom Kerberos configuration not passed to gen_tgt.sh

7.13.1

Cloudera Manager has a setting to specify a custom Kerberos configuration file location using the krb_krb5_conf_path parameter. However, this custom location is not passed to security/gen_tgt.sh from SecurityUtils::generateTgt. As a result, security/gen_tgt.sh defaults to /etc/krb5.conf, even when a custom path is configured.

To ensure the custom Kerberos configuration is used, set export KRB5_CONF IG=/*[***CUSTOM_PATH***]*/krb5/path in /etc/default/cloudera-scm-server to the same value as krb_krb5_conf_path.

OPSAPS-71669: The Continue option is disabled on the Static Service Pools Review page, affecting the functionality of Static Service Pools

7.13.1

7.13.1.100

The minimum and maximum I/O weight values for Cgroup v2 were incorrectly set to 100 and 1000, respectively, in Cloudera Manager 7.13.1.0. According to official Cgroup v2 documentation, the valid range should be 1 to 10,000. Due to this incorrect configuration range, the Continue option on the **Static Service Pools Review** page was disabled, preventing users from proceeding with pool configuration.

This issue might occur on clusters running Cloudera Manager 7.13.1.0 with Cgroup v2 resource management when configuring or reviewing Static Service Pools. After upgrading to Cloudera Manager 7.13.1.100 CHF-1, this issue no longer occurs.

None

OPSAPS-75290, OPSAPS-74994: The yarn_enable_container_usage_aggregation job is failing with “Null real user” error on Service Monitor.

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, and 7.13.1.500

The yarn_enable_container_usage_aggregation job is failing with "Null real user" error on Service Mnitor when the Yarn service is running on the computer cluster with Stub DFS, and when the Powerscale Service is running in the cluster with Powerscale DFS provider instead of HDFS.

None.

OPSAPS-71581: Cloudera Manager Agent's append_properties function fails with the realpath: invalid option -- 'u' error when executed from service control scripts.

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, and 7.13.1.500

Errors appear on the standard error (stderr) log of Cloudera Data Platform (CDP) services when you are attempting to trigger the cloudera-config.sh script. The error log contains the following message: realpath: invalid option -- 'u'. This is caused by an incorrectly placed command-line flag in the script, which prevents some service configurations from loading correctly.

To resolve this issue temporarily, you must perform the following workaround steps on each agent node in the base cluster::

1. Navigate to the directory /opt/cloudera/cm-agent/service/common/.
2. Open the cloudera-config.sh file for editing.
3. Locate the two lines that execute the python scripts such as append_properties.py and get_property.py.
4. In both lines, remove the -u flag or change its position to after python to the end of the line:

```
Change this:
value=$(python -u "${GET_PROPERTY_PY_DIR}"/get_property.py
"${1}" "${2}")
To this:
value=$(python "${GET_PROPERTY_PY_DIR}"/get_property.py "${1}"
 "${2}" -u)
```

```
Change this:
python -u "${APPEND_PROPERTIES_PY_DIR}"/append_properties.py
 "${1}" "${2}"
To this:
```

```
python "${APPEND_PROPERTIES_PY_DIR}"/append_properties.py "
${1}" "${2}" -u
```

5. After saving the changes on all agent nodes, restart the entire cluster for the new configuration to take effect.
6. Verify the fix by checking the stderr.log on a few service instances to ensure the realpath: invalid option -- 'u' error no longer appears.

OPSAPS-71878: Ozone fails to restart during cluster restart and displays the error message: Service has only 0 Storage Container Manager roles running instead of minimum required 1.

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, and 7.13.1.500

1. You must open Cloudera Manager on the second browser and restart the Ozone service separately.
2. After the Ozone service restarts, you can resume the cluster restart from the first browser.

ENGESC-30503, OPSAPS-74868: Cloudera Manager limited support for custom external repository requiring basic authentication

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, and 7.13.1.400

Current Cloudera Manager does not support custom external repository with basic authentication (the Cloudera Manager Wizard supports either HTTP (non-secured) repositories or usage of Cloudera https://archive.cloudera.com only). In case customers want to use a custom external repository with basic authentication, they might get errors.

The assumption is that you can access the external custom repository (such as Nexus or JFrog, or others) using LDAP credentials. In case an applicative user is used to fetch the external content (as done in Data Services with the docker imager repository), the customer should ensure that this applicative user is located under the user's base search path where the real users are being retrieved during LDAP authentication check (so the external repository will find it and will allow it to gain access for fetching the files).

Once done, you can use the current custom URL fields in the Cloudera Manager Wizard and enter the URL for the RPMs or parcels/other files in the format of "https://USERNAME:PASSWOR D@server.example.com/XX".

While using the password, you are advised to use only the printable ASCII character range (excluding space), whereas in case of a special character (not letter/number) it can be replaced with HEX value (For example, you can replace Aa1234$ with Aa1234%24 as '%24' is translated into $ sign).

OPSAPS-72164: Proxy Settings and Telemetry Publisher in Cloudera Manager

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, and 7.13.1.400

In Cloudera Manager 7.13.1, the PROXY settings for the Telemetry Publisher (TP) are not functioning as expected. This may impact the Telemetry Publisher's ability to communicate through a configured proxy.

You must upgrade to Cloudera Manager 7.13.1 CHF5 (7.13.1.500) and higher.

OPSAPS-60726: Newly saved parcel URLs are not showing up in the parcels page in the Cloudera Manager HA cluster.

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

To safely manage parcels in a Cloudera Manager HA environment, follow these steps:

1. Shutdown the Passive Cloudera Manager Server.
2. Add and manage the parcel as usual, as described in Install Parcels.
3. Restart the Passive Cloudera Manager server after parcel operations are complete.

OPSAPS-74341: NodeManagers might fail to start during the cluster restart after the Cloudera Manager 7.13.1.x upgrade

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

7.13.1.500

Cgroup v2 support is enabled in CDP 7.1.9 SP1 CHF5 and higher versions. However, if the user upgrades from Cloudera Manager 7.11.3.x to Cloudera Manager 7.13.1.x, and the environment is using cgroup v2, the NodeManagers might fail to start during the cluster restart after the Cloudera Manager 7.13.1.x upgrade.

To resolve this issue temporarily, you must perform the following steps:

1. Go to the YARN service page on the Cloudera Manager UI.
2. Navigate to the Configuration tab.
3. Search for NodeManager Advanced Configuration Snippet (Safety Valve) for yarn-site.xml.
4. Add the following entry:
   a. Add yarn.nodemanager.linux-container-executor.cgroups.v2.enabled=true
5. Restart the Nodemanagers. Nodemanagers restart successfully.

OPSAPS-73546: Service Monitor fails to perform Canary tests on HMS / HBASE / ZooKeeper due to missing dependencies

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

Due to a missing dependency caused by an incomplete build and packaging in certain OS releases, the HMS (Hive Metastore) Canary health test fails, logging a ClassNotFoundException in the Service Monitor log. This problem relates to all deliveries using runtime cluster version 7.1.x or 7.2.x, while the Cloudera Manager version is 7.13.1.x and the OS is NOT RHEL8.

In case your OS is either RHEL 9 or SLES 15 or Ubuntu 2004 or Ubuntu 2204 and if you install the Cloudera Manager 7.13.1.x version, then create a symbolic link using root user privileges on the node that host the Service Monitor service (cloudera-scm-firehose) at /opt/cloudera/cm/lib/cdh71/cdh71-hive-client-7.13.1-shaded.jar, pointing to /opt/cloudera/cm/lib/cdh7/cdh7-hive-client-7.13.1-shaded.jar.

Note: The above example relates to Cloudera Base on premises releases. In case your cluster is on Cloud, use "cdh72" instead of "cdh71" in the above symbolic link.

Restart the Service Monitor service post the change. This will allow the Service Monitor to perform Canary testing correctly on the HMS (Hive Metastore) service.

OPSAPS-72706, OPSAPS-73188: Hive queries fail after upgrading Cloudera Manager from 7.11.2 to 7.11.3 or later

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

7.13.1.500

Upgrading Cloudera Manager from version 7.11.2 or earlier to 7.11.3 or later causes Hive queries to fail due to JDK17 restrictions. Some JDK8 options are deprecated, leading to inaccessible classes and exceptions:

```
java.lang.reflect.InaccessibleObjectException: Unable to make fi
eld private volatile java.lang.String java.net.URI.string access
ible
```

To resolve this issue:

1. In Cloudera Manager, go to Tez Configuration
2. Append the following values to both tez.am.launch.cmd-opts and tez.task.launch.cmd-opts:

```
--add-opens=java.base/java.net=ALL-UNNAMED
--add-opens=java.base/java.util=ALL-UNNAMED
```

```
--add-opens=java.base/java.util.concurrent.atomic=ALL-UNNAMED
--add-opens=java.base/java.util.regex=ALL-UNNAMED
--add-opens=java.base/java.lang=ALL-UNNAMED
--add-opens=java.base/java.time=ALL-UNNAMED
--add-opens=java.base/java.io=ALL-UNNAMED
--add-opens=java.base/java.nio=ALL-UNNAMED
```

3. Save and restart

OPSAPS-72998: Missing charts for HMS event APIs

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

Charts for HMS event APIs (get_next_notification, get_current_notificationEventId, and fire_listener_event) are missing in Cloudera Manager Hive Hive Metastore Instance Charts Library API

Monitor HMS event activity using Hive Metastore logs.

OPSAPS-72270: Start ECS command fails on uncordon nodes step

7.13.1, 7.13.1.100, 7.13.1.200

7.13.1.300

In an ECS HA cluster, the server node restarts during the start up. This may cause the uncordon step to fail.

To resolve this issue temporarily, you must perform the following steps:

1. Run the following command on the same node to verify whether the kube-apiserver is ready:

```
kubectl get pods -n kube-system | grep kube-apiserver
```

2. Resume the command from the Cloudera Manager UI.

OPSAPS-73225: Cloudera Manager Agent reporting inactive/failed processes in Heartbeat request

7.13.1, 7.13.1.100, 7.13.1.200

7.13.1.300

As part of introducing Cloudera Manager 7.13.x, some changes were done to the Cloudera Manager logging, eventually causing Cloudera Manager Agent to report on inactive/stale processes during Heartbeat request.

As a result, the Cloudera Manager servers logs are getting filled rapidly with these notifications though they do not have impact on service.

In addition, with adding the support for the Cloudera Observability feature, some additional messages were added to the logging of the server. However, in case the customer did not purchase the Cloudera Observability feature, or the telemetry monitoring is not being used, these messages (which appears as "TELEMETRY_ALTUS_ACCOUNT is not configured for Otelcol" are filling the server logs and preventing proper follow-up on the server activities).

This will be fixed in a later release by moving these log notifications to DEBUG level so they don't appear on the Cloudera Manager server logs. Until that fix, perform the following workaround to filter out these messages.

On each of the Cloudera Manager servers, update with root credentials the file /etc/cloudera-scm-server/log4j.properties and add the following lines at the end of the file:

```
# === Custom Appender with Filters ===
log4j.appender.filteredlog=org.apache.log4j.ConsoleAppender
log4j.appender.filteredlog.layout=org.apache.log4j.PatternLayout
log4j.appender.filteredlog.layout.ConversionPattern=%d{ISO8601}
%p %c: %m%n
# === Filter #1: Drop warning ===
```

```
log4j.appender.filteredlog.filter.1=org.apache.log4j.varia.Str
ingMatchFilter
log4j.appender.filteredlog.filter.1.StringToMatch=Received Proces
s Heartbeat for unknown (or duplicate) process.
log4j.appender.filteredlog.filter.1.AcceptOnMatch=false
# === Filter #2: Drop telemetry config warning ===
log4j.appender.filteredlog.filter.2=org.apache.log4j.varia.String
MatchFilter
log4j.appender.filteredlog.filter.2.StringToMatch=TELEMETRY_ALTU
S_ACCOUNT is not configured for Otelcol
log4j.appender.filteredlog.filter.2.AcceptOnMatch=false
# === Accept all other messages ===
log4j.appender.filteredlog.filter.3=org.apache.log4j.varia.Accep
tAllFilter
# === Specific logger for AgentProtocolImpl ===
log4j.logger.com.cloudera.server.cmf.AgentProtocolImpl=WARN, filt
eredlog
log4j.additivity.com.cloudera.server.cmf.AgentProtocolImpl=false
# === Specific logger for BaseMonitorConfigsEvaluator ===
log4j.logger.com.cloudera.cmf.service.config.BaseMonitorConfigsE
valuator=WARN, filteredlog
log4j.additivity.com.cloudera.cmf.service.config.BaseMonitorConfi
gsEvaluator=false
```

Once done, restart the Cloudera Manager server(s) for the updated configuration to be picked.

OPSAPS-73211: Cloudera Manager 7.13.1 does not clean up Python Path impacting Hue to start

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.13.1 or higher versions with CDP Private Cloud Base 7.1.7.x Hue does not start because Cloudera Manager forces Hue to start with Python 3.8, and Hue needs Python 2.7.

The reason for this issue is because Cloudera Manager does not clean up the Python Path at any time, so when Hue tries to start the Python Path points to 3.8, which is not supported in CDP Private Cloud Base 7.1.7.x version by Hue.

To resolve this issue temporarily, you must perform the following steps:

1. Locate the hue.sh in /opt/cloudera/cm-agent/service/hue/.
2. Add the following line after export HADOOP_CONF_DIR=$CONF_DIR/hadoop-conf:

```
export PYTHONPATH=/opt/cloudera/parcels/CDH/lib/hue/build/env/
lib64/python2.7/site-packages
```

OPSAPS-73011: Wrong parameter in the /etc/default/cloudera-scm-server file

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300

7.13.1.400

In case the Cloudera Manager needs to be installed in High Availability (2 nodes or more as explained here), the parameter CMF_SERVER_ARGS in the /etc/default/cloudera-scm-server file is missing the word "export" before it (on the file there is only CMF_SERVER_ARGS= and not expo rt CMF_SERVER_ARGS=), so the parameter cannot be utilized correctly.

Edit the /etc/default/cloudera-scm-server file with root credentials and add the word "export" before the parameter CMF_SERVER_ARGS=.

OPSAPS-60346: Upgrading Cloudera Manager Agent triggers cert rotation in Auto-TLS use case 1

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

Upgrading Cloudera Manager Agent nodes from the Cloudera Manager UI wizard as part of a Cloudera Manager upgrade causes the host to get new certificates, which becomes disruptive.

The issue happens with use case 1 and Cloudera Manager DB is because Cloudera Manager always regenerates the host cert as part of the host install or host upgrade step. However, with use case 3, Cloudera Manager does not regenerate the cert as it comes from the user.

Currently, there are three following possible workarounds:

- Rotate all CMCA certs again using the `generateCmca` API command, and using the "location" argument to specify a directory on disk. This will revert to the old behavior of storing the certs on disk instead of the DB.

  

  Important: This approach is not recommended if Cloudera Manager is in HA mode or you plan to enable HA in the future.

- Switch to Auto-TLS Use Case 3 (Customer CA-signed Certificates).
- Manual upgrade of Cloudera Manager Agents, instead of upgrading from Cloudera Manager GUI.

OPSAPS-72447, CDPD-76705: Ozone incremental replication fails to copy renamed directory

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300

7.13.1.400

Ozone incremental replication using Ozone replication policies succeed but might fail to sync nested renames for FSO buckets.

When a directory and its contents are renamed between the replication runs, the outer level rename synced but did not sync the contents with the previous name.

None

OPSAPS-72756:The runOzoneCommand API endpoint fails during the Ozone replication policy run

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300

7.13.1.400

The `/clusters/{clusterName}/runOzoneCommand` Cloudera Manager API endpoint fails when the API is called with the `getOzoneBucketInfo` command. In this scenario, the Ozone replication policy runs also fail if the following conditions are true:

- The source Cloudera Manager version is 7.11.3 CHF11 or 7.11.3 CHF12.
- The target Cloudera Manager is version 7.11.3 through 7.11.3 CHF10 or 7.13.0.0 or later where the feature flag API_OZONE_REPLICATION_USING_PROXY_USER is disabled.

Choose one of the following methods as a workaround:

- Upgrade the target Cloudera Manager before you upgrade the source Cloudera Manager for 7.11.3 CHF12 version only.
- Pause all replication policies, upgrade source Cloudera Manager, upgrade destination Cloudera Manager, and unpause the replication policies.
- Upgrade source Cloudera Manager, upgrade target Cloudera Manager, and rerun the failed Ozone replication policies between the source and target clusters.

OPSAPS-65377: Cloudera Manager - Host Inspector not finding Psycopg2 on Ubuntu 20 or Redhat 8.x when Psycopg2 version 2.9.3 is installed.

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300

7.13.1.400

Host Inspector fails with Psycopg2 version error while upgrading to Cloudera Manager 7.13.1.x versions. When you run the Host Inspector, you get an error Not finding Psycopg2, even though it is installed on all hosts.

None

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

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

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

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

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-72804: For recurring replication policies, the interval is overwritten to 1 after the replication policy is edited

7.13.1

7.13.1.100

When you edit an Atlas, Iceberg, Ozone, or a Ranger replication policy that has a recurring schedule on the Replication Manager UI, the Edit Replication Policy modal window appears as expected. However, the frequency of the policy is reset to run at “1” unit where the unit depends on what you have set in the replication policy. For example, if you have set the replication policy to run every four hours, it is reset to one hour when you edit the replication policy.

After you edit the replication policy as required, you must ensure that you manually set the frequency to the original scheduled frequency, and then save the replication policy.

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

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

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

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

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-53160: Incorrect job run status appears for subsequent Hive ACID replication policy runs after the replication policy fails

7.13.1, 7.13.1.100, 7.13.1.200

7.13.1.300

When a Hive ACID replication policy run fails with the **FAILED_ADMIN** status, the subsequent Hive ACID replication policy runs show **SKIPPED** instead of **FAILED_ADMIN** status on the Cloudera Manager Replication Manager Replication Policies Actions Show History page which is incorrect. It is recommended that you check Hive ACID replication policy runs if multiple subsequent policy runs show the **SKIPPED** status.

None

CDPQE-36126: Iceberg replication fails when source and target clusters use different nameservice names

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

When you run an Iceberg replication policy between clusters where the source and target clusters use different nameservice names, the replication policy fails.

Perform the following steps to mitigate the issue, note that in the following steps the source nameservice is assumed to be ns1 and target cluster nameservice is assumed to be ns2:

1. Go to the Cloudera Manager Replication Replication Replication Policies page.
2. Click Actions Edit for the required Iceberg replication policy.
3. Go to the **Advanced** tab on the **Edit Iceberg Replication Policy** modal window.
4. Enter the mapreduce.job.hdfs-servers.token-renewal.exclude = ns1, ns2 key value pair for Advanced Configuration Snippet (Safety Valve) for source hdfs-site.xml and Advanced Configuration Snippet (Safety Valve) for destination hdfs-site.xml fields.
5. Save the changes.
6. Click Actions Run Now to run the replication policy.

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

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300

7.13.1.400

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

DMX-3973: Ozone replication policy with linked bucket as destination fails intermittently

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

When you create an Ozone replication policy using a linked/non-linked source cluster bucket and a linked target bucket, the replication policy fails during the "Trigger a OZONE replication job on one of the available OZONE roles" step.

None

OPSAPS-68143:Ozone replication policy fails for empty source OBS bucket

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

An Ozone incremental replication policy for an OBS bucket fails during the “Run File Listing on Peer cluster” step when the source bucket is empty.

None

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

7.13.1

7.13.1.100

During the Ozone replication policy run, Replication Manager does not read the value in the ozon e_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.

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

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-72204: HMS compaction configuration not updated through Cloudera Manager UI

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

The hive.compactor.initiator.on checkbox in Cloudera Manager UI for Hive Metastore (HMS) does not reflect the actual configuration value in cloud deployments. The default value is false, causing the compactor to not run.

To update the hive.compactor.initiator.on value:

1. In the Cloudera Manager, go to Hive Configuration
2. Add the value for hive.compactor.initiator.on to true in the "Hive Service Advanced Configuration Snippet (Safety Valve) for hive-site.xml"
3. Save the changes and Restart.

Once applied, the compaction process will run as expected.

OPSAPS-70702: Ranger replication policies fail if the clusters do not use AutoTLS

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

Ranger replication policies fail during the Exporting services, policies and roles from Ranger r emote 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

7.13.1

7.13.1.100

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 Sni ppet (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 Val ve) for ozone-conf/ozone-site.xml.

Additionally, you must ensure that you do not update the property values in Ozone Replication Ad vanced 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

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

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 7.1.8.x versions.

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

7.13.1

7.13.1.100

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

7.13.1

7.13.1.100

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

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

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

7.13.1

7.13.1.100

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.

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

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

7.13.1

7.13.1.100

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

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

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

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

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

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

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

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

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

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

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

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

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

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

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

7.13.1

7.13.1.100

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

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

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

None

OPSAPS-73138, OPSAPS-72435: Ozone OBS-to-OBS replication policies create directories in the target cluster even when no such directories exist on the source cluster

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

Ozone OBS-to-OBS replication uses Hadoop S3A connector to access data on the OBS buckets. Depending on the runtime version and settings in the clusters:

- directory marker keys (associated to the parent directories) appear in the destination bucket even when it is not available in the source bucket.
- delete requests of non-existing keys to the destination storage are submitted which result in `Key delete failed` messages to appear in the Ozone Manager log.

The OBS buckets are flat namespaces with independent keys, and the character '/' has no special significance in the key names. Whereas in FSO buckets, each bucket is a hierarchical namespace with filesystem-like semantics, where the '/' separated components become the path in the hierarchy. The S3A connector provides filesystem-like semantics over object stores where the connector mimics the directory behaviour, that is, it creates and optionally deletes the "empty directory markers". These markers get created when the S3A connector creates an empty directory. Depending on the runtime (S3A connector) version and settings, these markers are deleted when a descendant path is created and is not deleted.

Empty directory marker creation is inherent to S3A connector. Empty directory marker deletion behavior can be adjusted using the `fs.s3a.directory.marker.retention` = keep or delete key-value pair. For information about configuring the key-value pair, see Controlling the S3A Directory Marker Behavior.

OPSAPS-73655: Cloud replication fails after the delegation token is issued

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400

7.13.1.500

HDFS and Hive external table replication policies from an on-premises cluster to cloud fail when the following conditions are true:

1. You choose the Advanced Options Delete Policy Delete Permanently option during the replication policy creation process.

2. Incremental replication is in progress, that is the source paths of the replication are snapshottable directories and the bootstrap replication run is complete.

None

OPSAPS-75090: Ozone replication policies fail without source proxy user

7.13.1, 7.13.1.100, 7.13.1.200, 7.13.1.300, 7.13.1.400, 7.13.1.500

An Ozone replication policy with an empty Run on Peer as Username field (The default value for this field is empty) fails with the "java.io.IOException: Error acquiring writer for listing file "ofs://<service id>/user/om/.cm/distcp-staging/<timestamp>/fileList.seq": bucket name 'om' is too short, valid length is 3-63 characters". error message.

If you do not have a source proxy user name to specify in the Run on Peer as Username field, you can enter om as the default user for the replication on the source cluster.

Behavioral Changes

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.

Cloudera Manager 7.13.1 CHF5 (7.13.1.500)

Summary: Hive SYSDB creation failures

Previous behavior:

Any failure during the `Create SYSDB` command in Hive was treated as a fatal error, which would stop the entire Cloudera installation.

New behavior:

Now, Cloudera Manager treats failures during the creation of SYSDB in Hive as non-fatal. This allows the Cloudera installation to proceed even if the command fails. If users find that SYSDB does not exist, they can still create or recreate it manually from the **Actions** menu of the Hive Metastore Server role.

Cloudera Manager 7.13.1 CHF4 (7.13.1.400)

There are no behavioral changes in this release.

Cloudera Manager 7.13.1 CHF3 (7.13.1.300)

There are no behavioral changes in this release.

Cloudera Manager 7.13.1 CHF2 (7.13.1.200)

There are no behavioral changes in this release.

Cloudera Manager 7.13.1 CHF1 (7.13.1.100)

There are no behavioral changes in this release.

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 Base on premises and Cloudera on cloud version 7.3.1.

Fixed Common Vulnerabilities and Exposures in Cloudera Manager 7.13.1 and its cumulative hotfixes

Common Vulnerabilities and Exposures (CVE) that are fixed in Cloudera Manager 7.13.1 and its cumulative hotfixes.

Cloudera Manager 7.13.1 CHF5 (7.13.1.500)

CVEs	Package Name
CVE-2024-36114	Aircompressor
CVE-2019-10172	Jackson-mapper-asl
CVE-2024-22201	Jetty
CVE-2024-47554	Commons-io
CVE-2024-25710	Commons-compress
CVE-2024-26308	Commons-compress
CVE-2025-31672	Apache POI
CVE-2024-38820	Spring Framework
CVE-2025-22228	Spring Security

Cloudera Manager 7.13.1 CHF4 (7.13.1.400)

CVEs	Package Name
CVE-2025-27636	Apache Camel

Cloudera Manager 7.13.1 CHF3 (7.13.1.300)

CVEs	Package Name
CVE-2023-34442	Apache Camel
CVE-2023-24998	Apache Commons FileUpload
CVE-2024-53990	AsyncHttpClient

Cloudera Manager 7.13.1 CHF2 (7.13.1.200)

CVEs	Package Name
CVE-2025-23184	Apache CXF
CVE-2023-48795	sshj
CVE-2024-47072	XStream
CVE-2022-21699	IPython
CVE-2015-5607	IPython
CVE-2014-3429	IPython
CVE-2015-4707	IPython
CVE-2024-3651	Idna
CVE-2012-5783	httpclient
CVE-2014-3577	httpclient
CVE-2020-13956	httpclient
CVE-2015-5262	httpclient
CVE-2012-6153	httpclient

Cloudera Manager 7.13.1 CHF1 (7.13.1.100)

CVEs	Package Name
CVE-2023-44487	Netty
CVE-2024-21634	Ion-Java
CVE-2017-7536	Hibernate-Validator
CVE-2018-1000873	Jackson-databind
CVE-2017-15095	Jackson-databind
CVE-2017-17485	Jackson-databind
CVE-2017-7525	Jackson-databind
CVE-2018-11307	Jackson-databind
CVE-2018-14718	Jackson-databind
CVE-2018-14719	Jackson-databind
CVE-2018-7489	Jackson-databind
CVE-2019-14379	Jackson-databind

CVEs	Package Name
CVE-2019-14540	Jackson-databind
CVE-2019-14892	Jackson-databind
CVE-2019-16335	Jackson-databind
CVE-2019-16942	Jackson-databind
CVE-2019-16943	Jackson-databind
CVE-2019-17267	Jackson-databind
CVE-2019-17531	Jackson-databind
CVE-2019-20330	Jackson-databind
CVE-2020-8840	Jackson-databind
CVE-2020-9547	Jackson-databind
CVE-2020-9548	Jackson-databind
CVE-2020-10673	Jackson-databind
CVE-2018-5968	Jackson-databind
CVE-2020-10650	Jackson-databind
CVE-2020-24616	Jackson-databind
CVE-2020-24750	Jackson-databind
CVE-2020-35490	Jackson-databind
CVE-2020-35491	Jackson-databind
CVE-2020-36179	Jackson-databind
CVE-2020-36180	Jackson-databind
CVE-2020-36181	Jackson-databind
CVE-2020-36182	Jackson-databind
CVE-2020-36183	Jackson-databind
CVE-2020-36184	Jackson-databind
CVE-2020-36185	Jackson-databind
CVE-2020-36186	Jackson-databind
CVE-2020-36187	Jackson-databind
CVE-2020-36188	Jackson-databind
CVE-2020-36189	Jackson-databind
CVE-2021-20190	Jackson-databind
CVE-2018-12022	Jackson-databind
CVE-2019-12086	Jackson-databind
CVE-2019-14439	Jackson-databind
CVE-2020-36518	Jackson-databind
CVE-2022-42003	Jackson-databind
CVE-2022-42004	Jackson-databind
CVE-2019-12384	Jackson-databind
CVE-2019-12814	Jackson-databind
CVE-2020-13949	Libthrift

CVEs	Package Name
CVE-2018-1320	Libthrift
CVE-2019-0205	Libthrift
CVE-2019-0210	Libthrift
CVE-2018-11798	Libthrift
CVE-2024-38808	Spring Framework
CVE-2024-38829	Spring ldap
CVE-2024-38821	Spring Security
CVE-2024-38809	Spring Framework
CVE-2024-38816	Spring Framework
CVE-2024-38819	Spring Framework
CVE-2024-38820	Spring Framework

Cloudera Manager 7.13.1.0

CVEs	Package Name
CVE-2024-37891	urllib3
CVE-2023-43804	urllib3
CVE-2021-33503	urllib3
CVE-2020-26137	urllib3
CVE-2019-14893	Jackson-databind
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

CVEs	Package Name
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
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

CVEs	Package Name
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
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

CVEs	Package Name
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
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.

Deprecation Notices for Cloudera Manager

Certain features and functionality are deprecated or removed in Cloudera Manager 7.13.1. 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.

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