Cloudera Manager 7.11.3

Release Notes

Date published: 2020-11-30
Date modified: 2024-07-19

CLOUDERA

https://docs.cloudera.com/

Legal Notice

Contents

Cloudera Manager 7.11.3 Release Notes

Known issues, fixed issues and new features for Cloudera Manager 7.11.3 and CDP Private Cloud Base 7.1.9.

Important: Cloudera Manager 7.11.3 is nearing its end of life. Cloudera requests that all customers begin upgrading to Cloudera Manager 7.13.1. This release of Cloudera Manager, like previous releases, is backwards compatible with older versions of the Cloudera runtime. For Cloudera Manager 7.13.1 upgrade instructions, see Upgrading Cloudera Manager 7 to begin planning your upgrade, or see the Cloudera Support Lifecycle Policy page for more information.

Important:

Ubuntu 22.04, SLES 15 (SP4 and SP5) support will not be available for Cloudera Manager 7.11.3 CHF16. If you are using either the Ubuntu 22 or SLES 15 operating system, then do not update from Cloudera Manager 7.11.3 CHF15 to Cloudera Manager 7.11.3 CHF16.

To proceed with updating to CDP 7.1.9 SP1 CHF9, you must update from Cloudera Manager 7.11.3 CHF15 to Cloudera Manager 7.13.1 CHF4 (7.13.1.400). A Prerequisite to updating to Cloudera Manager 7.13.1 CHF4 is installing Python 3.11 on all hosts.

What's New in Cloudera Manager 7.11.3

New features and changed behavior for Cloudera Manager 7.11.3.

New features

Zero Downtime Upgrades (ZDU)

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

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

TLS 1.2 encryption support for secured database connections

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

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

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

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

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

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

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

TCPS support for connections to Oracle database

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

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

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

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

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

If you are running Cloudera Runtime 7.1.7 SP2 or below versions with Cloudera Manager 7.11.3, then Python 2.7 is still required for the Cloudera Runtime components. In this scenario, you must install both Python 2.7 (for Cloudera Runtime components) and Python 3 (for Cloudera Manager 7.11.3).

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

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

Support for noexec option on the /tmp directory

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

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

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

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

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

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

Certify CM with HA Postgres databases with SSL enabled

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

Iceberg replication policies

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

For more information, see Iceberg replication policies

Ranger replication policies

You can create Ranger replication policies in CDP Private Cloud Base Replication Manager. The Ranger replication policies migrate Ranger policies for HDFS, Hive, and HBase services between

Kerberos-enabled CDP Private Cloud Base 7.1.9 or higher clusters using Cloudera Manager 7.11.3 or higher versions.

For more information, see Ranger replication policies.

Incremental replication of Ozone data using Ozone replication policies

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

For more information, see Ozone replication policies.

Ozone snapshot policies

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

For more information, see Ozone snapshot policies.

Collecting Heartbeat data from Cloudera Manager

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

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

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

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

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

Changed or updated features

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

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

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

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

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

From this release onwards, you must import a new GPG public key into the OS key ring when installing the Cloudera Manager Agent, Cloudera Manager Server, and Cloudera Manager Daemon packages. The SHA-256 based signing key is applicable to both the fresh installation and upgrade

of Cloudera Manager (7.11.3 version). The new GPG keys are now signed with a more secure SHA-256 hashing algorithm.

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

Note:

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

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

Platform support for Cloudera Manager 7.11.3

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

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

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

What's new in Platform Support

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

Platform Support Enhancements

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

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

Known Issues in Cloudera Manager 7.11.3

Known issues in Cloudera Manager 7.11.3.

OPSAPS-75899: HDFS directory creation fails on JDK 11 or higher when LDAP or Active Directory integrated clusters using the hadoop.security.group.mapping property.

Due to this issue, many critical Cloudera Manager operations fail to complete, such as:

- Install Oozie ShareLib (Oozie Actions Install Oozie ShareLib)
- Install YARN MapReduce Framework JARs (YARN Actions Install YARN MapReduce Framework JARs)

You must perform the following workaround steps to manually add specific Java modules to the HDFS service script on all nodes in the cluster:

1. Navigate to the directory /opt/cloudera/cm-agent/service/hdfs/.
2. Open the hdfs.sh file for editing.
3. Locate the line starting with MKDIR_JAVA_OPTS.
4. Update it to include the following flags:

```
Change this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=f
ile://${MKDIR_LOG4J_FILE}"
To this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=fi
le://${MKDIR_LOG4J_FILE} --add-exports=java.naming/com.sun.j
ndi.ldap=ALL-UNNAMED --add-opens=java.naming/com.sun.jndi.ld
ap=ALL-UNNAMED"
```

OPSAPS-70915: Cloudera Manager Agent incorrectly detected its own cgroup path

The Cloudera Manager Agent incorrectly detected its own cgroup path and created the YARN NodeManager's cgroup (for example, hadoop-yarn) under the Cloudera Manager Agent's system.slice hierarchy instead of the root-level cgroup path.

For example, the YARN cgroup appeared as /sys/fs/cgroup/cpu,cpuacct/system.slice/cloudera-scm-agent.service/hadoop-yarn instead of /sys/fs/cgroup/cpu,cpuacct/hadoop-yarn.

Because of this misplacement, when you restart the Cloudera Manager Agent process, systemd automatically destroys the nested cgroup (/system.slice/cloudera-scm-agent.service/hadoop-yarn). This immediately kills all running YARN containers and causes active jobs to fail.

To prevent YARN from inheriting the Cloudera Manager Agent's cgroup hierarchy, you can explicitly configure Cloudera Manager Agent to use the root cgroup path. Perform this configuration by uncommenting and setting the cgroups paths in the Cloudera Manager Agent configuration file: /etc/cloudera-scm-agent/config.ini. Perform the following steps:

1. Under the [cgroups] section, uncomment or add the following lines (adjusting for your cgroup version and controller types):

```
[cgroups]
mounts=cpu,cpuacct,cpuset,memory
cpu_cgroup_mount_point=/sys/fs/cgroup/cpu,cpuacct
memory_cgroup_mount_point=/sys/fs/cgroup/memory
```

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

```
sudo systemctl restart cloudera-scm-agent
```

This configuration ensures that the Cloudera Manager Agent and its managed roles (such as YARN NodeManager) always use root-level cgroup paths rather than inheriting them from syst em.slice, and prevents systemd from automatically cleaning up those cgroups when Cloudera Manager Agent restarts.

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

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.

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

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

In Cloudera Manager 7.11.3, 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 7.11.3 CHF16 or higher.

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

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-73211: Cloudera Manager 7.11.3 does not clean up Python Path impacting Hue to start

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.11.3 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

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-72984: Alerts due to change in hostname fetching functionality in jdk 8 and jdk 11

Upgrading JAVA from JDK 8 to JDK 11 creates the following alert in CMS:

Bad : CMSERVER:pit666.slayer.mayank: Reaching Cloudera Manager Server failed

This happens due to a functionality change in JDK 11 on hostname fetching.

```
[root@pit666.slayer ~]# /us/lib/jvm/java-1.8.0/bin/java GetHostN
ame
Hostname: pit666.slayer.mayank

[root@pit666.slayer ~]# /usr/lib/jvm/java-11/bin/java GetHostName
Hostname: pit666.slayer
```

You can notice that the "hostname" is set to a short name instead of FQDN.

The current workaround is to set the hostname as FQDN.

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 7.11.3 or Cloudera Manager 7.11.3 CHF-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-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.

Manually add = sign to file_metadata_reload_properties configuration and modify the Gflags configuration file when the file_metadata_reload_properties configuration is passed as empty.

OPSAPS-60169: Cloudera Manager statestore connectivity health check fails if Kerberos is enabled for Impala Web UI

If Kerberos is enabled for Impala Web UI, the Cloudera Manager statestore connectivity health check fails and the Service Monitor displays the following exception:

```
WARN com.cloudera.cmon.firehose.polling.CdhTask: (14 skipped) Ex
ception in doWork for task: impala_IMPALA_SERVICE_STATE_FETCHER
```

In Cloudera Manager, go to Clusters Impala Configuration , search for the "Enable Kerberos Authentication for HTTP Web-Consoles" property and disable this property.

For more information, see the Cloudera Knowledge Base article.

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

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

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

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

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

OPSAPS-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.

Upgrade to CDP 7.1.9, or downgrade to Cloudera Manager 7.7.1.

OPSAPS-70207: Cloudera Manager Agents sending the Impala profile data with an incorrect header

Cloudera Manager agent might send incorrect HTTP header to Telemetry Publisher causing incorrect Content-Type error message resulting connection error. This issue causes missing Impala profile on Cloudera Observability.

Impala profile data is not available on Cloudera Observability.

Telemetry Publisher logs show:

DEBUG org.apache.cxf.jaxrs.utils.JAXRSUtils: No method match, method name : addProfileEvent, request path : /cluster/impala2, method @Path : /{clusterName}/{serviceName}, HTTP Method : POST, method HTTP Method : POST, ContentType : application/x-www-form-urlencoded, method @Consumes : application/json,, Accept : */*,, method @Produces : application/json,.

Cloudera Manager agent logs on Impalad hosts report:

Error occurred when sending entry to server: HTTP Error 415: Unsupported Media Type, url: http://<telemetry_publisher_host>:<port>

None

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

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

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

OPSAPS-68845: Cloudera Manager Server fails to start after the Cloudera Manager upgrade

Starting from the Cloudera Manager 7.11.3 version up to the Cloudera Manager 7.11.3 CHF7 version, the Cloudera Manager Server fails to start after the Cloudera Manager upgrade due to Navigator user roles improperly handled in the upgrade in some scenarios.

None

OPSAPS-68577: Invalid Iceberg license validator message

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

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

You can safely suppress the Configuration validator message.

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

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

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

None

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

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

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

Azul Open JDK 8

RHEL or SLES

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

Ubuntu or Debian

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

Azul Open JDK 11

For DEBs only

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

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

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

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

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

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

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

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

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

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

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

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

Use CLI prompts instead of GUI prompts.

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

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

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

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

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

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

You must restart the active and passive Cloudera Manager nodes.

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

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

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

1. Stop the Cloudera Manager Server.

```
sudo systemctl stop cloudera-scm-server
```

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

```
select * from CONFIGS WHERE ATTR='keystore_type';
```

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

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

6. Start the Cloudera Manager Server.

```
sudo systemctl start cloudera-scm-server
```

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

You must reach out to Cloudera Support.

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

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

- MariaDB 10.6.15
- MariaDB 10.5.22
- MariaDB 10.4.31

None

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

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

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

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

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

OPSAPS-68484: Hive queries fail with 'get_partitions_ps_with_auth_req' error

Hive queries fail with the error due to a mismatch between HiveServer2 and Hive metastore during a zero-downtime upgrade (ZDU).

```
Error: Invalid method name: 'get_partitions_ps_with_auth_req'
```

The issue is addressed by adjusting the upgrade order, ensuring HiveMetastore is upgraded before HiveServer2.

DMX-3167

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

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

DMX-3193

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

Copy the metadata file from the target cluster to the source cluster and run the incremental replication again.

OPSAPS-68143

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

DMX-3169

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

Provide a proxy user to submit the DistCp jobs.

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

DMX-3174

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

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

mapreduce.job.hdfs-servers.token-renewal.exclude = *[***SOURCE NAME SERVICE***], [***TARGET NAME SERVICE***]*.

For more information, see Kerberos setup guidelines.

CDPD-59437

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

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

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

1. You must uninstall Cloudera Manager Agent package manually.

2. Install Python 3 on the host before upgrading to Cloudera Manager 7.11.3. See Installing Python 3.
3. Reinstall the Cloudera Manager Agent.

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

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

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

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

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

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

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

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

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

DMX-3003

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

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

DMX-2977, DMX-2978

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

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

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

When Cloudera Manager is upgraded to 7.11.3, if the CDP cluster is not 7.1.9, then the YARN Container Usage Aggregation job fails.

Add the following property in the MapReduce Client Advanced Configuration Snippet (Safety Valve) for mapred-site.xml file.

```
NAME: mapreduce.jvm.add-opens-as-default
VALUE: false
```

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

On a cluster with a custom krb5.conf file location configured in Cloudera Manager, HDFS HTTPFS role is not able to start because it does not have the custom Kerberos configuration file setting properly propagated to the service, and therefore it fails with a Kerberos related exception: in thread "main" java.io.IOException: Unable to initialize WebAppContext at org.apache.hadoop.http.HttpServer2.start(HttpServer2.java:1240) at org.apache.hadoop.fs.http.server.HttpFSServerWebServer.start(HttpFSServerWebServer.java:131) at org.apache.hadoop.fs.http.server.HttpFSServerWebServer.main(HttpFSServerWebServer.java:162) Caused by: java.lang.IllegalArgumentException: Can't get Kerberos realm at org.apache.hadoop.security.HadoopKerberosName.setConfiguration(HadoopKerberosName.java:71) at org.apache.hadoop.security.UserGroupInformation.initialize(UserGroupInformation.java:329) at org.apache.hadoop.security.UserGroupInformation.setConfiguration(UserGroupInformation.java:380) at org.apache.hadoop.lib.service.hadoop.FileSystemAccessService.init(FileSystemAccessService.java:166) at org.apache.hadoop.lib.server.BaseService.init(BaseService.java:71) at org.apache.hadoop.lib.server.Server.initServices(Server.java:581) at org.apache.hadoop.lib.server.Server.init(Server.java:377) at org.apache.hadoop.fs.http.server.HttpFSServerWebApp.init(HttpFSServerWebApp.java:100) at org.apache.hadoop.lib.servlet.ServerWebApp.contextInitialized(ServerWebApp.java:158) at org.eclipse.jetty.server.handler.ContextHandler.callContextInitialized(ContextHandler.java:1073) at org.eclipse.jetty.servlet.ServletContextHandler.callContextInitialized(ServletContextHandler.java:572) at org.eclipse.jetty.server.handler.ContextHandler.contextInitialized(ContextHandler.java:1002) at org.eclipse.jetty.servlet.ServletHandler.initialize(ServletHandler.java:765) at org.eclipse.jetty.servlet.ServletContextHandler.startContext(ServletContextHandler.java:379) at org.eclipse.jetty.webapp.WebAppContext.startWebapp(WebAppContext.java:1449) at org.eclipse.jetty.webapp.WebAppContext.startContext(WebAppContext.java:1414) at org.eclipse.jetty.server.handler.ContextHandler.doStart(ContextHandler.java:916) at org.eclipse.jetty.servlet.ServletContextHandler.doStart(ServletContextHandler.java:288) at org.eclipse.jetty.webapp.WebAppContext.doStart(WebAppContext.java:524) at org.eclipse.jetty.util.component.AbstractLifeCycle.start(AbstractLifeCycle.java:73) at org.eclipse.jetty.util.component.ContainerLifeCycle.start(ContainerLifeCycle.java:169) at org.eclipse.jetty.util.component.ContainerLifeCycle.doStart(ContainerLifeCycle.java:117) at org.eclipse.jetty.server.handler.AbstractHandler.doStart(AbstractHandler.java:97) at org.eclipse.jetty.util.component.AbstractLifeCycle.start(AbstractLifeCycle.java:73) at org.eclipse.jetty.util.component.ContainerLifeCycle.start(ContainerLifeCycle.java:169) at org.eclipse.jetty.server.Server.start(Server.java:423) at org.eclipse.jetty.util.component.ContainerLifeCycle.doStart(ContainerLifeCycle.java:110) at org.eclipse.jetty.server.handler.AbstractHandler.doStart(AbstractHandler.java:97) at org.eclipse.jetty.server.Server.doStart(Server.java:387) at org.eclipse.jetty.util.component.AbstractLifeCycle.start(AbstractLifeCycle.java:73) at org.apache.hadoop.http.HttpServer2.start(HttpServer2.java:1218) ... 2 more Caused by: java.lang.IllegalArgumentException: KrbException: Cannot locate default realm at java.security.jgss/javax.security.auth.kerberos.KerberosPrincipal.<init>(KerberosPrincipal.java:174) at org.apache.hadoop.security.authentication.util.KerberosUtil.getDefaultRealm(KerberosUtil.java:108) at org.apache.hadoop.security.HadoopKerberosName.setConfiguration(HadoopKerberosName.java:69) ...

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

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

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

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

Upgrade the source Cloudera Manager to 7.11.3 or higher version.

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

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

Contact Cloudera support for a workaround.

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

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

None.

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

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

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

CDPD-75871: Table column stats replication fails on HA cluster, with multi threaded import

After you upgrade, the Hive external table replication policies fail during the “Hive import” step if the following conditions are true:

- The default setting columnStatsImportMultiThreaded = true in the Cloudera Manager Clusters Hive service Configuration hive_replication_env_safety_valve property is retained after the upgrade process.
- You are using HA Hive Metastore, and are using Hive and Impala to query source tables.

Perform the following steps to resolve the issue:

1. Go to the target Cloudera Manager Clusters Hive service Configuration tab.
2. Locate the hive_replication_env_safety_valve property.
3. Add COLUMN_STATS_IMPORT_MULTI_THREADED=false, and Save the changes.
4. Restart the Hive service.
5. Run the Hive external table replication policy.

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

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

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

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

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

A replicated table is dropped automatically in the target cluster when the following conditions are true:

- You create an Iceberg replication policy to replicate T1 (database D1 has tables T1 and T2).
- You drop T1 in the source cluster.
- You edit the replication policy to remove T1 in the include table regex pattern, and add T2.
- You run the replication policy.
- T2 is replicated and T1 is dropped in the target cluster.

In such scenarios, it is recommended not to drop the tables on the source cluster.

None

Fixed Issues in Cloudera Manager 7.11.3

Fixed issues in Cloudera Manager 7.11.3.

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

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

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

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

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

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

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

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

This issue is fixed now.

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

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

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

OPSAPS-67942: Installation failed due to schematool error

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

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

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

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

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

OPSAPS-63724

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

Activating this workaround:

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

Usage notes:

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

```
INFO util.DistCpUtils: There are changes on target, falling
back to regular distcp
NFO distcp.PreCopyListingCheck: The changes on target are safe
 to ignore.
```

```
INFO distcp.PreCopyListingCheck: Note that it is up to the
downstream processing steps if it falls back to full file li
sting or continue with snapshot diff execution
INFO distcp.PreCopyListingCheck: Changes to target: true
INFO distcp.PreCopyListingCheck: Changes to target are safe to
 ignore: true
```

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

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

Allowed changes:

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

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

OPSAPS-63930

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

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

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

Activating this workaround:

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

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

Usage notes: There can be conflicting source side renames and rename - delete interactions when their destination side replay need to use temporary renames (for example, a name swap between two paths using three renames). For these cases, the temporary rename destination will typically be next

to the final rename destination (will share the same parent path) avoiding both above mentioned failure scenarios. Such temporary renames will be logged during execution like:

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

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

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

OPSAPS-66197

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

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

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

Workaround:

- Set the HDFS service core-site.xml advanced configuration snippet (also called safety valve) (on the destination side) "com.cloudera.enterprise.distcp.overwrite-merge-existing-rename-targets.enabled" to "true". (Note that enabling the workaround in OPSAPS-63724 uses a different advanced configuration snippet).
- In an incremental replication run, check the stderr log of the last "Trigger a HDFS replication job on one of the available HDFS roles." step and make sure that the INFO distcp.DistCpSync: Overwrite merge of already existing move targets is enabled message is displayed.

Usage notes:

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

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

OPSAPS-63558

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

The following new counters are added:

- FILES_MOVED_TO_COMMON_TEMP_DIR: Number of files and directories moved to a common temporary directory to be renamed or deleted later in the process.

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

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

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

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

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

- FFL replication schedules all the missing paths for deletion regardless of parent relationships. When the `com.cloudera.enterprise.distcp.parent-only-delete.enabled` safety valve is set to "true", only the topmost deleted paths are scheduled for deletions and their descendants or children cannot be accessed. This is optional and by default turned off (which preserves the previous behavior).
- Delete requests can be issued from multiple threads concurrently to improve performance, and can be enabled and configured using the following safety valves:
  - `com.cloudera.enterprise.distcp.parallel-ffl-delete.enabled`. Default is "false".
  - `com.cloudera.enterprise.distcp.parallel-ffl-delete.threads`. Default is 20.
  - `com.cloudera.enterprise.distcp.parallel-ffl-delete.max-queue-size`. Default is 10000.

OPSAPS-65823

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

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

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

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

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

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

OPSAPS-65104

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

To resolve this issue, perform the following steps:

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

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

     This ensures that the column statistics import is skipped entirely.

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

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

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

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

OPSAPS-65267

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

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

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

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

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

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

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

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

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

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

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

Cloudera Manager displays the following error message:

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

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

OPSAPS-67897, OPSAPS-68023

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

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

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

Cloudera Manager 7.11.3 CHF17

- CVE-2024-38827 - VMware
- CVE-2024-38820 - Spring Framework
- CVE-2025-22228 - Spring Security

Cloudera Manager 7.11.3 CHF16

- CVE-2023-34442 - Apache Camel
- CVE-2025-27636 - Apache Camel

Cloudera Manager 7.11.3 CHF15

No CVEs are fixed in this release.

Cloudera Manager 7.11.3 CHF14

- CVE-2023-48795 - sshj

Cloudera Manager 7.11.3 CHF13

- CVE-2024-47072 - XStream

Cloudera Manager 7.11.3 CHF12

- CVE-2024-25710 - Commons-Compress
- CVE-2024-26308 - Commons-Compress
- CVE-2017-7536 - Hibernate-Validator

- CVE-2024-36114 - Aircompressor
- CVE-2020-13949 - libthrift
- CVE-2024-38829 - Spring Framework
- CVE-2024-38821 - Spring Security
- CVE-2024-38816 - Spring Framework
- CVE-2024-38819 - Spring Framework
- CVE-2024-38820 - Spring Framework

Cloudera Manager 7.11.3 CHF11

- CVE-2024-38808 - Spring Framework
- CVE-2024-38809 - Spring Framework

Cloudera Manager 7.11.3 CHF10

- CVE-2024-37891 - urllib3
- CVE-2023-43804 - urllib3
- CVE-2021-33503 - urllib3
- CVE-2020-26137 - urllib3
- CVE-2023-0833 - Okhttp
- CVE-2021-0341 - Okhttp
- CVE-2023-3635 - Okio

Cloudera Manager 7.11.3 CHF9.1

- CVE-2024-39689 - Certifi
- CVE-2023-33202 - Bouncy Castle
- CVE-2024-34447 - Bouncy Castle
- CVE-2024-29857 - Bouncy Castle
- CVE-2024-30171 - Bouncy Castle
- CVE-2023-33201 - Bouncy Castle
- CVE-2024-29736 - Apache CXF
- CVE-2024-32007 - Apache CXF

Cloudera Manager 7.11.3 CHF8

- CVE-2021-28170 - Javax.el
- CVE-2023-36478 - Eclipse Jetty
- CVE-2023-40167 - Eclipse Jetty
- CVE-2023-36479 - Eclipse Jetty
- CVE-2023-41900 - Eclipse Jetty
- CVE-2023-52428 - Nimbus-jose-jwt
- CVE-2023-20862 - Spring Security
- CVE-2024-22257 - Spring Security
- CVE-2023-20859 - Spring Vault
- CVE-2024-22262 - Spring Web

Cloudera Manager 7.11.3 CHF7

- CVE-2023-26048 - Eclipse Jetty
- CVE-2023-26049 - Eclipse Jetty
- CVE-2023-39196 - Apache Ozone
- CVE-2024-1597 - Postgresql

- CVE-2022-1471 - Snakeyaml
- CVE-2024-23944 - Apache Zookpeer

Cloudera Manager 7.11.3 CHF6

- 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-2021-46877 - Jackson Databind
- CVE-2020-13697 - Nanohttpd
- CVE-2022-21230 - Nanohttpd
- CVE-2024-22243 - Spring Framework

Cloudera Manager 7.11.3 CHF5

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

Cloudera Manager 7.11.3 CHF4

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

Cloudera Manager 7.11.3 CHF3

- CVE-2022-46364 - Apache CXF
- CVE-2022-46363 - Apache CXF
- CVE-2022-1415 - Drools
- CVE-2021-41411 - Drools
- CVE-2022-41853 - Hsqldb

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

Cloudera Manager 7.11.3 CHF2

- CVE-2022-25647 - Gson
- CVE-2021-28165 - Eclipse Jetty
- CVE-2022-2048 - Eclipse Jetty
- CVE-2020-27223 - Eclipse Jetty
- CVE-2021-28169 - Eclipse Jetty
- CVE-2021-34428 - Eclipse Jetty
- CVE-2021-28163 - Eclipse Jetty
- CVE-2022-2047 - Eclipse Jetty
- CVE-2022-45688 - org.json
- CVE-2023-5072 - org.json
- CVE-2023-3635 - Okio

Cloudera Manager 7.11.3 CHF1

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

Cloudera Manager 7.11.3

- CVE-2021-25642
- CVE-2022-25168
- CVE-2022-31129
- CVE-2021-36373
- CVE-2021-36374
- CVE-2020-9493
- CVE-2022-23305
- CVE-2023-26464

- CVE-2018-14721
- CVE-2018-14718
- CVE-2018-14719
- CVE-2018-14720
- CVE-2018-19360
- CVE-2018-19361
- CVE-2018-19362
- CVE-2018-12022
- CVE-2018-12023
- CVE-2022-36364
- CVE-2017-15095
- CVE-2018-5968
- CVE-2020-28491
- CVE-2022-40146
- CVE-2022-41704
- CVE-2022-42890
- CVE-2022-38398
- CVE-2022-38648
- CVE-2020-15522
- CVE-2020-0187
- CVE-2020-26939
- CVE-2020-13955
- CVE-2019-14862
- CVE-2023-24998
- CVE-2022-23457
- CVE-2022-24891
- CVE-2018-11792
- CVE-2021-28131
- CVE-2018-11785
- CVE-2022-21724
- CVE-2022-26520
- CVE-2022-31197
- CVE-2022-41946
- CVE-2021-27905
- CVE-2021-44548
- CVE-2021-29943
- CVE-2020-13941
- CVE-2017-3163
- CVE-2017-3164
- CVE-2018-1308
- CVE-2019-12401
- CVE-2019-0193
- CVE-2015-8795
- CVE-2015-8796
- CVE-2015-8797
- CVE-2018-11802
- CVE-2020-5421
- CVE-2022-22978
- CVE-2021-22112
- CVE-2022-22976

- CVE-2016-1000027
- CVE-2020-5397
- CVE-2022-40152
- CVE-2022-40151
- CVE-2022-41966
- CVE-2020-10683
- CVE-2018-1000632
- CVE-2014-0229
- CVE-2014-3627
- CVE-2014-3566
- CVE-2013-4221
- CVE-2013-4271
- CVE-2017-14868
- CVE-2017-14949
- CVE-2014-1868
- CVE-2018-8010
- CVE-2018-8026
- CVE-2017-5637
- CVE-2021-37533
- CVE-2018-14040
- CVE-2022-37865
- CVE-2022-37866
- CVE-2013-2035
- CVE-2014-125087
- CVE-2021-33813
- CVE-2014-3643
- CVE-2022-40149
- CVE-2022-40150
- CVE-2022-45685
- CVE-2022-45693
- CVE-2023-1436
- CVE-2017-7657
- CVE-2017-7658
- CVE-2017-7656
- CVE-2017-9735
- CVE-2020-27216
- CVE-2019-10247
- CVE-2019-10241
- CVE-2018-12536
- CVE-2019-10246
- CVE-2016-5725
- CVE-2023-1370
- CVE-2021-37714
- CVE-2022-36033
- CVE-2016-4000
- CVE-2018-1320
- CVE-2019-0205
- CVE-2019-0210
- CVE-2018-11798
- CVE-2019-17571

- CVE-2021-4104
- CVE-2016-2402
- CVE-2021-27807
- CVE-2021-27906
- CVE-2021-31811
- CVE-2021-31812
- CVE-2022-26336
- CVE-2022-1471
- CVE-2017-18640
- CVE-2022-25857
- CVE-2022-38749
- CVE-2022-38751
- CVE-2022-38752
- CVE-2022-41854
- CVE-2022-38750
- CVE-2017-8028
- CVE-2019-11272
- CVE-2019-3795
- CVE-2019-14887
- CVE-2022-23437
- CVE-2020-11988

Deprecation notices in Cloudera Manager 7.11.3

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

Terminology

Items in this section are designated as follows:

Deprecated

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

Moving

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

Removed

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

Deprecation Notices for Cloudera Manager

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

Removed

SHA-1-GNU Privacy Guard (GPG) keys

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

Platform and OS

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

Database Support:

The listed databases are deprecated from the 7.11.3 release.

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

Operating System

The listed operating system is removed from the 7.11.3 release.

- Ubuntu 18.04

Note: Ubuntu 18 Operating System is not supported from Cloudera Manager 7.11.3 upto Cloudera Manager 7.11.3 CHF6 versions. You must upgrade the Operating System from Ubuntu 18 to Ubuntu 20 before you upgrade to Cloudera Manager 7.11.3 CHF6. For performing major OS upgrade, see Upgrading the Operating System to a new Major Version.

Cloudera Manager 7.11.3 Cumulative hotfix 7 (CDP Private Cloud Base 7.1.9 SP1)

Known issues, fixed issues and new features for Cloudera Manager 7.11.3 CHF7 and CDP Private Cloud Base 7.1.9 SP1.

Important: Cloudera Manager 7.11.3 CHF7 and later versions support CDP Private Cloud Base 7.1.9 SP1.

What's New in Cloudera Manager 7.11.3 Cumulative hotfix 7 (CDP Private Cloud Base 7.1.9 SP1)

You must be aware of the major enhancements or changes and additions or corrections for the Cloudera Manager 7.11.3 Cumulative hotfix 7 release. Learn how the new improvements benefit you.

New features

Kerberos (MIT and AD) authentication support for MariaDB

Adding Kerberos (MIT and AD) authentication support for MariaDB on Cloudera Manager clusters (includes TLS 1.2 and non-TLS 1.2 clusters). Kerberos is a network authentication protocol that provides security for your cluster.

Now you can enable Kerberos authentication on MariaDB Database Server in the database environment. See Enabling Kerberos (MIT and AD) autentication for MariaDB Database Server.

Adding Python 3.10 support on Ubuntu 22.04, SLES 15 SP4, and SLES 15 SP5

Adding Python 3.10 support on Ubuntu 22.04, SLES 15 SP4, and SLES 15 SP5 from Cloudera Manager 7.11.3 CHF7 release.

Important:

Due to a change in support from Python 3.8 to Python 3.10 for SLES 15 SP4 and SLES 15 SP5, only a regular upgrade of Cloudera Manager to 7.11.3 CHF7 and CDP Runtime cluster to 7.1.9 SP1 is possible and must occur sequentially without starting the cluster between the Cloudera Manager and CDP Runtime cluster upgrades.

You must install Python 3.10 for SLES 15 SP4 and SLES 15 SP5 on all hosts before upgrading to Cloudera Manager 7.11.3 CHF7. See Installing Python 3.

Important: Cloudera Manager now requires Python 3.10 on all versions of SLES 15, including SLES 15 SP4. It is not possible to support two different versions of Python for the same major version of operating system. If the cluster was previously running Python 3.8, then you must upgrade to Python 3.10.

For more information about the operating systems that are supported when using Python 3.x with the Cloudera Manager Agents, see Python support for Cloudera Manager 7.11.3 CHF7.

Replicate Atlas metadata and data lineage using replication policies in Replication Manager (technical preview)

You can use one of the following methods to replicate Atlas metadata and data lineage for Hive external tables and Iceberg tables:

- Choose Replicate Atlas Metadata on the **General** tab during the Hive external table replication policy creation or edit process to replicate the metadata and data lineage associated with the chosen Hive external tables.
- Choose Replicate Atlas Metadata on the **General** tab during the Iceberg replication policy creation or edit process to replicate the metadata and data lineage associated with the chosen Iceberg tables.
- Create Atlas replication policy to replicate the metadata and data lineage of all the Hive external tables, Iceberg tables, and any other Atlas supported entities in the source cluster to the target cluster.

This is a technical preview feature. It is *not* recommended for production deployments. Cloudera recommends that you try this feature in development or test environments. To enable this feature, contact your Cloudera account team.

For more information, see Atlas replication policies.

Iceberg replication policy enhancements

During the Iceberg replication policy creation process you can:

- replicate Iceberg tables residing in a custom directory using the General Alternate target data root option.
- enter a username to run the MapReduce job other than the hdfs user using the General Run as Username option.
- choose to copy the table column statistics associated with the chosen Iceberg tables using the General Replicate Table Column Statistics option.

- configure the following options on the **Advanced** tab:
  - Choose Use Batch Size, and then enter the maximum number of snapshots to process for an export batch.
  - Add one or more key-value pairs for the following properties:
    - `Advanced Configuration Snippet (Safety Valve) for source hdfs-site.xml`
    - `Advanced Configuration Snippet (Safety Valve) for source core-site.xml`
    - `Advanced Configuration Snippet (Safety Valve) for destination hdfs-site.xml`
    - `Advanced Configuration Snippet (Safety Valve) for destination core-site.xml`

For more information, see Creating Iceberg replication policies.

Ozone snapshot policy enhancements

On the Cloudera Manager Clusters *OZONE SERVICE* Bucket Browser tab, you can

- view and browse the list of volumes and buckets in the service.
- view the list of snapshots for a bucket.
- create a snapshot

Additionally, you can restore or delete a snapshot.

For more information, see Restoring Ozone snapshots in Cloudera Manager

Hive external table replication policy enhancement (technical preview)

You can accomplish metadata-only replication for Ozone storage-backed Hive external tables after you specify a valid Destination Staging Path during the Hive external table replication policy creation process. You must replicate the data using Ozone replication policies.

This is a technical preview feature. It is not recommended for production deployments. Cloudera recommends that you try this feature in development or test environments. To enable this feature, contact your Cloudera account team.

For more information, see Creating Hive external table replication policy

Replication Manager support for AWS temporary credentials (technical preview)

You can use temporary AWS credentials, through the IDBroker service, to replicate HDFS data, Hive external tables, and HBase data from 7.1.9 SP1 Kerberized CDP Private Cloud Base clusters using Cloudera Manager 7.11.3 CHF7 or higher versions to S3 buckets or CDP Public Cloud clusters on AWS. You can also use the temporary AWS credentials to replicate HDFS data from S3 buckets to 7.1.9 SP1 Kerberized CDP Private Cloud Base clusters or higher using Cloudera Manager 7.11.3 CHF7 or higher versions.

This is a technical preview feature. It is not recommended for production deployments. Cloudera recommends that you try this feature in development or test environments. To enable this feature, contact your Cloudera account team.

For more information, see Add IDBroker to use temporary AWS session credentials

High availability support for Impala

Impala High availability is now configurable through Cloudera Manager 7.11.3 CHF7 or higher versions.

Custom properties atlas.jaas.KafkaClient.option.password was available in a clear text format in CDP cluster services when Kerberos authentication was not present.

To provide a secured access, two new fields are introduced for username / password and a radio field for loginModule for Kerberos or Plain selection.

atlas.jaas.KafkaClient.option.username=username

atlas.jaas.KafkaClient.option.password=<password is in clear text>

atlas.jaas.KafkaClient.option.loginModuleName=KERBEROS(default)

Changed or updated features

Deploy client configuration command timed-out on larger node clusters

The Deploy Client Config command is improved now. Previously, it could take long and time-out on large clusters. It is now leveraging multithreading and optimized for parallel execution. The command is now expected to complete much faster and should not cause timeouts.

The performance of `GenerateCredentials` on MIT is improved now

Multi-threaded credential generation for MIT has been enabled, allowing principal generation scripts to run in parallel.

Python support for Cloudera Manager 7.11.3 CHF7

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

Python 3.10
• Ubuntu 22.04 • SLES 15 SP4 • SLES 15 SP5 • RHEL 9.2 (LTS) • Oracle 9.2 UEK

What's new in Platform Support

You must be aware of the platform support for the Cloudera Manager 7.11.3 CHF7 release.

Platform Support Enhancements

- **New OS support**:
  - RHEL 9.4
  - RHEL 9.2
  - RHEL 8.10
  - RHEL 8.8 FIPS (extended RHEL 8.8 support for FIPS customers with JDK 11 and JDK 17 versions)
  - Oracle 8.10 (supported with Red Hat Compatible Kernel (RHCK) only)
  - Oracle 9.2
  - Oracle 9.2 UEK
  - Ubuntu 22.04
  - SLES 15 SP5

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

- **New Database support**:
  - PostgreSQL 16
  - PostgreSQL 15
  - MariaDB 10.11
  - Oracle 23c
  - Oracle 21c
  - Oracle 19c
- **New JDK Version**:
  - Support for FIPS + OpenJDK 17 (From Cloudera Runtime 7.1.9 SP1 release onwards). For more information, see JDK FIPS prerequisites for CDP.
  - Azul Open JDK 8 and Azul Open JDK 11 are supported with Cloudera Manager 7.11.3 CHF7 and higher versions.

Known Issues in Cloudera Manager 7.11.3 Cumulative hotfix 7 (CDP Private Cloud Base 7.1.9 SP1)

Known issues in Cloudera Manager 7.11.3 Cumulative hotfix 7.

OPSAPS-75899: HDFS directory creation fails on JDK 11 or higher when LDAP or Active Directory integrated clusters using the hadoop.security.group.mapping property.

Due to this issue, many critical Cloudera Manager operations fail to complete, such as:

- Install Oozie ShareLib (Oozie Actions Install Oozie ShareLib)
- Install YARN MapReduce Framework JARs (YARN Actions Install YARN MapReduce Framework JARs)

You must perform the following workaround steps to manually add specific Java modules to the HDFS service script on all nodes in the cluster:

1. Navigate to the directory /opt/cloudera/cm-agent/service/hdfs/.
2. Open the hdfs.sh file for editing.
3. Locate the line starting with MKDIR_JAVA_OPTS.
4. Update it to include the following flags:

```
Change this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=f
ile://${MKDIR_LOG4J_FILE}"
To this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=fi
le://${MKDIR_LOG4J_FILE} --add-exports=java.naming/com.sun.j
ndi.ldap=ALL-UNNAMED --add-opens=java.naming/com.sun.jndi.ld
ap=ALL-UNNAMED"
```

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

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_CONFIG=/*[***CUSTOM_PATH***]*/krb5/path in /etc/default/cloudera-scm-server to the same value as krb_krb5_conf_path.

OPSAPS-70915: Cloudera Manager Agent incorrectly detected its own cgroup path

The Cloudera Manager Agent incorrectly detected its own cgroup path and created the YARN NodeManager's cgroup (for example, hadoop-yarn) under the Cloudera Manager Agent's system.slice hierarchy instead of the root-level cgroup path.

For example, the YARN cgroup appeared as /sys/fs/cgroup/cpu,cpuacct/system.slice/cloudera-scm-agent.service/hadoop-yarn instead of /sys/fs/cgroup/cpu,cpuacct/hadoop-yarn.

Because of this misplacement, when you restart the Cloudera Manager Agent process, systemd automatically destroys the nested cgroup (/system.slice/cloudera-scm-agent.service/hadoop-yarn). This immediately kills all running YARN containers and causes active jobs to fail.

To prevent YARN from inheriting the Cloudera Manager Agent's cgroup hierarchy, you can explicitly configure Cloudera Manager Agent to use the root cgroup path. Perform this configuration by uncommenting and setting the cgroups paths in the Cloudera Manager Agent configuration file: /etc/cloudera-scm-agent/config.ini. Perform the following steps:

1. Under the [cgroups] section, uncomment or add the following lines (adjusting for your cgroup version and controller types):

```
[cgroups]
mounts=cpu,cpuacct,cpuset,memory
cpu_cgroup_mount_point=/sys/fs/cgroup/cpu,cpuacct
memory_cgroup_mount_point=/sys/fs/cgroup/memory
```

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

```
sudo systemctl restart cloudera-scm-agent
```

This configuration ensures that the Cloudera Manager Agent and its managed roles (such as YARN NodeManager) always use root-level cgroup paths rather than inheriting them from syst em.slice, and prevents systemd from automatically cleaning up those cgroups when Cloudera Manager Agent restarts.

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

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.

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

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

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-73211: Cloudera Manager 7.11.3 does not clean up Python Path impacting Hue to start

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.11.3 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

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-72984: Alerts due to change in hostname fetching functionality in jdk 8 and jdk 11

Upgrading JAVA from JDK 8 to JDK 11 creates the following alert in CMS:

Bad : CMSERVER:pit666.slayer.mayank: Reaching Cloudera Manager Server failed

This happens due to a functionality change in JDK 11 on hostname fetching.

```
[root@pit666.slayer ~]# /us/lib/jvm/java-1.8.0/bin/java GetHostN
ame
Hostname: pit666.slayer.mayank

[root@pit666.slayer ~]# /usr/lib/jvm/java-11/bin/java GetHostName
Hostname: pit666.slayer
```

You can notice that the "hostname" is set to a short name instead of FQDN.

The current workaround is to set the hostname as FQDN.

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 7.11.3 or Cloudera Manager 7.11.3 CHF-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-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.

Manually add = sign to file_metadata_reload_properties configuration and modify the Gflags configuration file when the file_metadata_reload_properties configuration is passed as empty.

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.

To resolve this issue, disable the Enable Log Event Capture configuration on the Configuration page of a service.

OPSAPS-68845: Cloudera Manager Server fails to start after the Cloudera Manager upgrade

Starting from the Cloudera Manager 7.11.3 version up to the Cloudera Manager 7.11.3 CHF7 version, the Cloudera Manager Server fails to start after the Cloudera Manager upgrade due to Navigator user roles improperly handled in the upgrade in some scenarios.

None

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

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

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

OPSAPS-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.

Upgrade to CDP 7.1.9, or downgrade to Cloudera Manager 7.7.1.

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

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

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

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

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

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

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

The status of the deleted table displays as ACTIVE.

None

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

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

None

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

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

Upgrade the source Cloudera Manager to 7.11.3 or higher version.

OPSAPS-70861: HDFS replication policy creation process fails for cluster with Isilon

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

Create the HDFS replication policy using Cloudera Manager REST APIs.

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

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

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

None

OPSAPS-70704: Kerberos connectivity check does not work as expected with JDK17 when you add Cloudera Manager peers

When you add a source Cloudera Manager that supports JDK17, the Kerberos connectivity check fails and the Error while reading /etc/krb5.conf on <hostname ; for all hosts>... error message appears.

None

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

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

None

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

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

None

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

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

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

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

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

DMX-3659: Updated row is replicated as inserted row

During an Iceberg replication policy run, an updated row on the source cluster is replicated as an inserted row which is incorrect. This results in a discrepancy because the number of rows for the table in the source and target clusters do not match.

None

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

During the Ozone replication policy run, Replication Manager does not read the value in the 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-71067: Wrong interval sent from the Replication Manager UI after Ozone replication policy submit or edit process.

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

OPSAPS-72509, CDPD-32440: Hive metadata transfer to GCS fails with ClassNotFoundException

Hive replication policies from an on-premises cluster to cloud fails during the “Transfer Metadata Files” step if the following conditions are true:

- the target is a GCS Data Lake
- 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 is because the fs.gs.delegation.token.binding property is already defined in the configuration and cannot be unset to disable the delegation tokens in the cloud connector service.

None

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

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

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

A replicated table is dropped automatically in the target cluster when the following conditions are true:

- You create an Iceberg replication policy to replicate T1 (database D1 has tables T1 and T2).
- You drop T1 in the source cluster.
- You edit the replication policy to remove T1 in the include table regex pattern, and add T2.
- You run the replication policy.
- T2 is replicated and T1 is dropped in the target cluster.

In such scenarios, it is recommended not to drop the tables on the source cluster.

None

Fixed Issues in Cloudera Manager 7.11.3 Cumulative hotfix 7 (CDP Private Cloud Base 7.1.9 SP1)

Fixed issues in Cloudera Manager 7.11.3 Cumulative hotfix 7.

OPSAPS-69018: Cloudera Manager fails to support multiple SAML role values

When multiple values for the SAML role assignment attribute are returned in an assertion, Cloudera Manager only reads the first attribute value returned in an assertion list.

Since the attribute typically reflects a user’s LDAP groups, multiple values are common and can include any number of values which may or may not be mapped to roles in Cloudera Manager, in

any order. This can cause authorization failures, or unexpected limited access rights in Cloudera Manager. This issue is fixed now.

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

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

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

OPSAPS-68387: Cloudera Manager UI incorrectly showing Skipped status for Hive ACID replication policy jobs when the job status was unknown

The **Status** column on the Cloudera Manager Replication Replication Policies page was incorrectly showing **Skipped** for Hive ACID replication policy jobs when the job status was unknown. The column now shows the **Waiting for Update** status for the Hive ACID replication policy jobs until the job status is confirmed.

OPSAPS-68494: Replication metric getter handles scenarios when "hive.resultset.use.unique.column.names" = "false"

Replication Metric getter failed when the `hive.resultset.use.unique.column.names` parameter was set to false because the resulting columns were non-unique. The Replication Metric getter now configures the `hive.resultset.use.unique.column.names` parameter to `true` during its JDBC session to override the service configuration.

OPSAPS-69978: Cruise Control capacity.py script fails on Python 3.8

Cruise Control no longer fails to start on Python 3.x when the capacity information is queried during the startup process. The script querying the capacity information is now fully compatible with Python 3.x.

OPSAPS-70269: Mismatch in ssl_enabled configuration between Cloudera Manager and Knox

The ssl.enabled property is populated in gateway-site.xml. Knox startup check script works when ssl is not enabled. This issue is fixed now.

OPSAPS-70257: Cloudera Manager upgrade fails with an error

While using CDP 7.1.6, if you try to upgrade Cloudera Manager with a version prior to 7.11.3 CHF7, the Cloudera Manager upgrade failed with the following error message:

```
Start RANGER_KMS-1 FAILED with Failed to start service
```

This issue is fixed now.

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-70051: Configuration issue with the hive.server2.tez.initialize.default.sessions parameter

Cloudera Manager incorrectly sets hive.server2.tez.initialize.default.sessions to true, conflicting with its expected false value in Hive configurations.

Adjusted Cloudera Manager to align with Hive configuration, ensuring the parameter defaults correctly to false for consistency and to prevent overriding settings.

OPSAPS-70248: Optimize Impala Graceful Shutdown Initiation Time

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

OPSAPS-68906: Impala Rolling Restart Sequence for ZDU

This adjustment refines the Impala rolling restart sequence with catalog and statestore HA support, reducing Impala downtime during cluster upgrades.

OPSAPS-67641: The Next Run column for Hive ACID replication policies shows the correct message.

The **Next Run** column on the Cloudera Manager Replication Replication Policies page showed **None Scheduled** for recurring Hive ACID replication policy jobs, which is incorrect. The column now displays the correct message.

OPSAPS-68246: Added a parameter for ozoneReplicationResult response for Ozone replication policies

The resultMessage parameter in the `ozoneReplicationResult` response for Ozone replication policies in Cloudera Manager REST API shows whether the replication command completed successfully or has failed with a specific message.

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-70607: Peer name validation step during Iceberg replication policy creation process is updated

During the Iceberg replication policy creation process if the source cluster name is renamed, the replication policy creation process does not fail.

OPSAPS-70492: ZDU | Handling of JDK add-opens flag in YARN with Cloudera Manager

The `JDK_JAVA_OPTIONS` environment variable is now used to supply the JDK 17 related flags.

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-69859: Correct configuration propagation in Cloudera Manager for non-HA clusters

In non-HA clusters, Cloudera Manager previously failed to propagate a few configurations for Ozone Manager (OM). This resulted in errors when attempting to submit a DistCp job to YARN, causing the submission process to fail. This issue has been fixed to ensure all required configurations are propagated correctly in non-HA clusters, allowing DistCp job submissions to proceed without errors.

OPSAPS-69987: Set the decommissioning state during decommission of Ozone Manager and Storage Container Manager

This issue is resolved by setting the state of master roles (OM and SCM) in Ozone to decommissioned after successful decommissioning.

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-63193: Need to enable Atlas canary check by default

Atlas canary check was disabled because Data Hub creation fails as Data Lake Atlas service health degrades.

OPSAPS-70355: Change compression from 'gz' to 'SNAPPY' in Atlas HBase tables

Changed the compression algorithm from GZ to SNAPPY in Atlas HBase tables to reduce the compaction time.

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-70299: Added optional Run As User option for hbase initial snapshot export on the source cluster

Added `runAsUser` query parameter to the `clusters/{cluster-name}/services/{service-name}/snapshots/hbase/remote` endpoint. When creating an on-premises to cloud HBase replication policies with Perform Initial Snapshot option, this appears as the Export snapshot user field in the **Create HBase replication policy** wizard in Cloudera Replication Manager.

When a user is specified in the `runAsUser` parameter, the YARN application that exports the HBase snapshot gets submitted by the *hbase* user impersonating the specified *runAsUser*.

To ensure the YARN application succeeds, the *hbase* user must be allowed to impersonate the *runAsUser* in HDFS by configuring the required properties and values in the HDFS core_site_safety_valve.

For example, if you want to allow the impersonation from any host and if the `runAsUser` is in the *repl* user group, you can set the following key-value pairs in Cloudera Manager Clusters *[***CORE SETTINGS***]* Configuration Cluster-wide Advanced Configuration Snippet (Safety Valve) for core-site.xml :

- hadoop.proxyuser.hbase.groups = repl
- hadoop.proxyuser.hbase.hosts = *

OPSAPS-68704: Admin server failed to start when the modified “iceberg-replicaton-batch.jar” is deployed

The admin server failed to start when you deployed the modified iceberg-replicaton-batch.jar. The runtime dependencies were removed from the modified JAR file to reduce its file size as these dependencies were provided by the runtime components. This issue is resolved.

Fixed Common Vulnerabilities and Exposures in Cloudera Manager 7.11.3 cumulative hotfix 7

Common Vulnerabilities and Exposures (CVE) that are fixed in Cloudera Manager 7.11.3 cumulative hotfix 7.

Cloudera Manager 7.11.3 CHF7

- CVE-2023-26048 - Eclipse Jetty
- CVE-2023-26049 - Eclipse Jetty
- CVE-2023-39196 - Apache Ozone
- CVE-2024-1597 - Postgresql
- CVE-2022-1471 - Snakeyaml
- CVE-2024-23944 - Apache Zookpeer

Deprecation notices in Cloudera Manager 7.11.3 Cumulative hotfix 7

Certain features and functionalities have been removed or deprecated in Cloudera Manager 7.11.3 Cumulative hotfix 7. You must review these items to understand whether you must modify your existing configuration. You can also learn about the features that will be removed or deprecated in the future release to plan for the required changes.

Terminology

Items in this section are designated as follows:

Deprecated

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

Moving

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

Removed

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

Platform and OS

The listed Operating Systems and databases are deprecated or removed from the Cloudera Manager 7.11.3 Cumulative hotfix 7 release.

Database Support:

The following databases are deprecated from the Cloudera Manager 7.11.3 CHF7 release:

- PostgreSQL 12
- MariaDB 10.4
- MySQL 5.7

The following databases are removed and no longer supported from the Cloudera Manager 7.11.3 CHF7 release:

- PostgreSQL 11

Operating System

The following operating systems are deprecated from the Cloudera Manager 7.11.3 CHF7 release:

- RHEL 8.6
- RHEL 7.9
- RHEL 7.9 (FIPS)

- CentOS 7.9
- SLES 12 SP5

Cumulative hotfixes

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

Important: Cloudera Manager 7.11.3 is nearing its end of life. Cloudera requests that all customers begin upgrading to Cloudera Manager 7.13.1. This release of Cloudera Manager, like previous releases, is backwards compatible with older versions of the Cloudera runtime. For Cloudera Manager 7.13.1 upgrade instructions, see Upgrading Cloudera Manager 7 to begin planning your upgrade, or see the Cloudera Support Lifecycle Policy page for more information.

Cloudera Manager 7.11.3 Cumulative hotfix 17

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 17.

This cumulative hotfix was released on August 13, 2025.

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

New features and changed behavior for Cloudera Manager 7.11.3 CHF17 (version: 7.11.3.39-69579823): Rocky Linux 9.4 support for Cloudera Manager

Starting with the Cloudera Manager 7.11.3 CHF17 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 7.11.3 CHF17 release.

Upgraded embedded PostgreSQL to 14.16

The embedded PostgreSQL version within Key Trustee Server is upgraded from 14.2 to 14.16.

OPSAPS-74756 and OPSAPS-74460: Previously, Spark extractions did not fetch YARN application metadata. The Spark jobs could not fetch accurate queue information and did not produce an auxiliary-files/YARN/appInfo.json file in the extraction output.

With the new functionality, Spark extractions now include YARN application metadata. This provides an accurate queue mapping for Spark jobs and creates an auxiliary-files/YARN/appInfo.json file in the extraction output.

The new configuration properties, extractor.spark.yarn.app.max.retry and extractor.spark.yarn.app.retry.wait.millis, control the retry attempts and wait time for YARN API calls. By default, the system attempts a YARN API call only once. However, you can configure these properties for more retries to increase resilience using Cloudera Manager.

1. Log in to Cloudera Manager
2. Navigate to Cloudera Management Service
3. Click Configuration
4. Search for Telemetry Publisher Advanced Configuration Snippet (Safety Valve) for telemetrypublisher.conf
5. Add the following parameters extractor.spark.yarn.app.max.retry and extractor.spark.yarn.app.retry.wait.millis

   For example,

   ```
   extractor.spark.yarn.app.max.retry=5
   ```

```
extractor.spark.yarn.app.retry.wait.millis=1000
```

OPSAPS-74300: Allow override of the Cloudera Manager supplied PYTHONPATH in Livy CSDs

Livy uses the Python executable and PYTHONPATH as set by the Cloudera Manager Agent for PySpark sessions. If required, now it is possible to override these default settings via multiple environment variables.

You can override the default settings for PYTHONPATH using the following methods:

Livy

1. In Cloudera Manager Clusters Livy Configuration Override CM PYTHONPATH enabled check the checkbox.
2. In Cloudera Manager Clusters Livy Configuration Override CM PYTHONPATH value enter the new **Livy Server Default Group**.

Livy for Spark 3

1. In Cloudera Manager Clusters LIVY_FOR_SPARK3 Configuration Override CM PYTHONPATH enabled check the **Livy Server for Spark 3 Default Group** checkbox.
2. In Cloudera Manager Clusters LIVY_FOR_SPARK3 Configuration Override CM PYTHONPATH value enter the **Livy Server for Spark 3 Default Group** value.

Note: For more information, see Setting Python path variables for Livy in Cloudera Runtime 7.1.9.

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF17 (version: 7.11.3.39-69579823):
OPSAPS-75899: HDFS directory creation fails on JDK 11 or higher when LDAP or Active Directory integrated clusters using the hadoop.security.group.mapping property.

Due to this issue, many critical Cloudera Manager operations fail to complete, such as:

- Install Oozie ShareLib (Oozie Actions Install Oozie ShareLib)
- Install YARN MapReduce Framework JARs (YARN Actions Install YARN MapReduce Framework JARs)

You must perform the following workaround steps to manually add specific Java modules to the HDFS service script on all nodes in the cluster:

1. Navigate to the directory /opt/cloudera/cm-agent/service/hdfs/.
2. Open the hdfs.sh file for editing.
3. Locate the line starting with MKDIR_JAVA_OPTS.
4. Update it to include the following flags:

```
Change this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=f
ile://${MKDIR_LOG4J_FILE}"
To this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=fi
le://${MKDIR_LOG4J_FILE} --add-exports=java.naming/com.sun.j
ndi.ldap=ALL-UNNAMED --add-opens=java.naming/com.sun.jndi.ld
ap=ALL-UNNAMED"
```

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

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.

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

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

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-73211: Cloudera Manager 7.11.3 does not clean up Python Path impacting Hue to start

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.11.3 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-72984: Alerts due to change in hostname fetching functionality in jdk 8 and jdk 11

Upgrading JAVA from JDK 8 to JDK 11 creates the following alert in CMS:

Bad : CMSERVER:pit666.slayer.mayank: Reaching Cloudera Manager Server failed

This happens due to a functionality change in JDK 11 on hostname fetching.

```
[root@pit666.slayer ~]# /us/lib/jvm/java-1.8.0/bin/java GetHostN
ame
Hostname: pit666.slayer.mayank

[root@pit666.slayer ~]# /usr/lib/jvm/java-11/bin/java GetHostName
Hostname: pit666.slayer
```

You can notice that the "hostname" is set to a short name instead of FQDN.

The current workaround is to set the hostname as FQDN.

OPSAPS-72784: Upgrades from CDH6 to CDP Private Cloud Base 7.1.9 SP1 or higher versions fail with a health check timeout exception

If you are using Cloudera Manager 7.11.3 cumulative hotfix 14 or higher versions and upgrading from CDH 6 to CDP Private Cloud Base 7.1.9 SP1 or higher versions, the upgrade fails with a CMUpgradeHealthException timeout exception. This is because upgrades from CDH 6 to CDP Private Cloud Base 7.1.9 SP1 or to any of its cumulative hotfix versions are not supported.

None.

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

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

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

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

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

Ensure that both the instances use the same Java version. If it is not possible to have the same Java versions on both the instances, ensure that they use the same encryption type for Kerberos. To check

the encryption type in Cloudera Manager, search for krb_enc_types on the Cloudera Manager Administration Settings page.

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

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

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

OPSAPS-70771: Running replication policy runs must not allow you to download the performance reports

During a replication policy run, the A server error has occurred. See Cloudera Manager server log for details error message appears on the UI and the Cloudera Manager log shows "java.lang.IllegalStateException: Command has no result data." when you click:

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

This is because the Replication Manager UI shows the performance report links as enabled and clickable which is incorrect. You can download the reports only after the replication job run is complete.

None

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

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

None

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

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

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-74398: Ozone and HDFS replication policies might fail when you use different destination proxy user and source proxy user

HDFS on-premises to on-premises replication fails when the following conditions are true:

- You configure different Run As Username and Run on Peer as Username during the replication policy creation process.
- The user configured in Run As Username does not have the permission to access the source path on the source HDFS.

Ozone replication fails when the following conditions are true:

- FSO-to-FSO replication or an OBS-to-OBS replication with Incremental with fallback to full file listing or Incremental only replication type.

- You configured different Run As Username and Run on Peer as Username during the replication policy creation process.
- The user configured in Run As Username does not have the permission to access the source bucket on the source Ozone.

Provide the same permissions to the user configured in Run As Username as the permissions of Run on Peer as Username on the source cluster.

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

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.

Following are the list of fixed issues that were shipped for Cloudera Manager 7.11.3 CHF17 (version: 7.11.3.39-69579823):

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

This issue is fixed now. Health port added as a configuration to the Knox configuration. The health topology port can be set with topology port mapping and by setting the new configuration the chec kDeployment script will use the new health port.

OPSAPS-73711 and OPSAPS-73165: When Ranger is enabled, Telemetry Publisher fails to export Hive payloads from Data Hub due to the missing Ranger client dependencies in the Telemetry Publisher classpath.

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

OPSAPS-74379 and OPSAPS-74375: 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-72439, OPSAPS-74265: HDFS and Hive external tables replication policies failed when using custom krb5.conf files

The issue appeared because the custom krb5.conf was not propagated to the required files. To mitigate this issue, complete the instructions provided in Step 13 in Using a custom Kerberos configuration path before you run the replication polices.

OPSAPS-73602, OPSAPS-74360: HDFS replication policies to cloud failed with HTTP 400 error

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

OPSAPS-74040, OPSAPS-74057: 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.

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

OPSAPS-73655, OPSAPS-74060: Cloud replication failed after the delegation token was issued

When you chose the Advanced Setting Delete Policy Delete permanently option during the replication policy creation process, the HDFS and Hive external table replication policies from an on-premises cluster to cloud failed when incremental replication was in progress (the source paths

of the replication were snapshottable directories and the bootstrap replication run was complete). This issue is fixed.

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

OPSAPS-73645, OPSAPS-73846: 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-70403: Custom Kerberos configuration not passed to gen_tgt.sh

Previously, the custom Kerberos configuration file location specified by the krb_krb5_conf_path parameter was not passed to security/gen_tgt.sh from SecurityUtils::generateTgt. As a result, security/gen_tgt.sh defaulted to /etc/krb5.conf, even when a custom path was configured.

This issue is now resolved. security/gen_tgt.sh correctly uses the custom Kerberos configuration file specified by the krb_krb5_conf_path parameter, which means that if a custom path is configured then security/gen_tgt.sh will use the custom */[***CUSTOM_PATH***]*/etc/krb5.conf file. If no custom path is set, security/gen_tgt.sh defaults to using the /etc/krb5.conf file.

Fixed Common Vulnerabilities and Exposures

For information about Common Vulnerabilities and Exposures (CVE) that are fixed in Cloudera Manager 7.11.3 cumulative hotfix 17, see Fixed Common Vulnerabilities and Exposures in Cloudera Manager 7.11.3 cumulative hotfixes.

The repositories for Cloudera Manager 7.11.3 CHF17 are listed in the following table:

Table 1: Cloudera Manager 7.11.3 CHF17

Repository Type	Repository Location

RHEL 9 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.39/redhat9/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.39/redhat9/yum/cloudera-manager.repo`

RHEL 8 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.39/redhat8/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.39/redhat8/yum/cloudera-manager.repo`

Repository Type	Repository Location

RHEL 7 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.39/redhat7/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.39/redhat7/yum/cloudera-manager.repo`

SLES 15	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.39/sles15/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.39/sles15/yum/cloudera-manager.repo`

SLES 12	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.39/sles12/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.39/sles12/yum/cloudera-manager.repo`

Ubuntu 22	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.39/ubuntu2204/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.39/ubuntu2204/apt/cloudera-manager.list`

Ubuntu 20	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.39/ubuntu2004/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.39/ubuntu2004/apt/cloudera-manager.list`

Cloudera Manager 7.11.3 Cumulative hotfix 16

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 16.

This cumulative hotfix was released on July 15, 2025.

Important:

Ubuntu 22.04, SLES 15 (SP4 and SP5) support will not be available for Cloudera Manager 7.11.3 CHF16. If you are using either the Ubuntu 22 or SLES 15 operating system, then do not update from Cloudera Manager 7.11.3 CHF15 to Cloudera Manager 7.11.3 CHF16.

To proceed with updating to CDP 7.1.9 SP1 CHF9, you must update from Cloudera Manager 7.11.3 CHF15 to Cloudera Manager 7.13.1 CHF4 (7.13.1.400). A Prerequisite to updating to Cloudera Manager 7.13.1 CHF4 is installing Python 3.11 on all hosts.

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

New features and changed behavior for Cloudera Manager 7.11.3 CHF16 (version: 7.11.3.36-67636814):
Cgroup v2 support on RHEL 9 for Cloudera Manager 7.11.3 CHF16

Important: Support for Cgroup v2 will start from Cloudera Manager 7.11.3 CHF16 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. 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 7.11.3 CHF16, must disable Cgroup v2 if already enabled before upgrading to Cloudera Manager 7.11.3 CHF16 version as cgroup v2 is not supported with Cloudera Manager version lower than 7.11.3 CHF16.
- 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.

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

Trino plugin support in Ranger has been added.

OPSAPS-70457: 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.

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF16 (version: 7.11.3.36-67636814):
OPSAPS-75899: HDFS directory creation fails on JDK 11 or higher when LDAP or Active Directory integrated clusters using the hadoop.security.group.mapping property.

Due to this issue, many critical Cloudera Manager operations fail to complete, such as:

- Install Oozie ShareLib (Oozie Actions Install Oozie ShareLib)
- Install YARN MapReduce Framework JARs (YARN Actions Install YARN MapReduce Framework JARs)

You must perform the following workaround steps to manually add specific Java modules to the HDFS service script on all nodes in the cluster:

1. Navigate to the directory /opt/cloudera/cm-agent/service/hdfs/.
2. Open the hdfs.sh file for editing.
3. Locate the line starting with MKDIR_JAVA_OPTS.
4. Update it to include the following flags:

```
Change this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=f
ile://${MKDIR_LOG4J_FILE}"
To this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=fi
le://${MKDIR_LOG4J_FILE} --add-exports=java.naming/com.sun.j
ndi.ldap=ALL-UNNAMED --add-opens=java.naming/com.sun.jndi.ld
ap=ALL-UNNAMED"
```

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

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.

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

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-74288: Alert publisher cannot send email alerts due to missing JAR

Alert publisher cannot send email alerts due to missing camel-attachments-3.14.9.jar in Cloudera Manager.

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

1. SSH login into Cloudera Manager cluster.
2. Manually add the camel-attachments-3.14.9.jar file into Cloudera Manager at /opt/cloudera/cm/lib/ directory.
3. Restart the Cloudera Manager server by running the following command:

```
sudo systemctl restart cloudera-scm-server
```

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

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-73211: Cloudera Manager 7.11.3 does not clean up Python Path impacting Hue to start

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.11.3 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-72984: Alerts due to change in hostname fetching functionality in jdk 8 and jdk 11

Upgrading JAVA from JDK 8 to JDK 11 creates the following alert in CMS:

Bad : CMSERVER:pit666.slayer.mayank: Reaching Cloudera Manager Server failed

This happens due to a functionality change in JDK 11 on hostname fetching.

```
[root@pit666.slayer ~]# /us/lib/jvm/java-1.8.0/bin/java GetHostN
ame
Hostname: pit666.slayer.mayank

[root@pit666.slayer ~]# /usr/lib/jvm/java-11/bin/java GetHostName
Hostname: pit666.slayer
```

You can notice that the "hostname" is set to a short name instead of FQDN.

The current workaround is to set the hostname as FQDN.

OPSAPS-72784: Upgrades from CDH6 to CDP Private Cloud Base 7.1.9 SP1 or higher versions fail with a health check timeout exception

If you are using Cloudera Manager 7.11.3 cumulative hotfix 14 or higher versions and upgrading from CDH 6 to CDP Private Cloud Base 7.1.9 SP1 or higher versions, the upgrade fails with a CMUpgradeHealthException timeout exception. This is because upgrades from CDH 6 to CDP Private Cloud Base 7.1.9 SP1 or to any of its cumulative hotfix versions are not supported.

None.

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

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

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

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

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

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

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

MariaDB 10.6, by default, includes the property require_secure_transport=ON in the configuration file (/etc/my.cnf), which is absent in MariaDB 10.4. This setting prohibits non-TLS connections,

leading to access issues. This problem is observed in High Availability (HA) mode, where certain operations may not be using the same connection.

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

OPSAPS-70771: Running replication policy runs must not allow you to download the performance reports

During a replication policy run, the A server error has occurred. See Cloudera Manager server log for details error message appears on the UI and the Cloudera Manager log shows "java.lang.IllegalStateException: Command has no result data." when you click:

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

This is because the Replication Manager UI shows the performance report links as enabled and clickable which is incorrect. You can download the reports only after the replication job run is complete.

None

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

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

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

None

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

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

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-74398: Ozone and HDFS replication policies might fail when you use different destination proxy user and source proxy user

HDFS on-premises to on-premises replication fails when the following conditions are true:

- You configure different Run As Username and Run on Peer as Username during the replication policy creation process.
- The user configured in Run As Username does not have the permission to access the source path on the source HDFS.

Ozone replication fails when the following conditions are true:

- FSO-to-FSO replication or an OBS-to-OBS replication with Incremental with fallback to full file listing or Incremental only replication type.
- You configured different Run As Username and Run on Peer as Username during the replication policy creation process.

- The user configured in Run As Username does not have the permission to access the source bucket on the source Ozone.

Provide the same permissions to the user configured in Run As Username as the permissions of Run on Peer as Username on the source cluster.

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

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.

Following are the list of fixed issues that were shipped for Cloudera Manager 7.11.3 CHF16 (version: 7.11.3.36-67636814):

OPSAPS-73585, OPSAPS-73432: 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-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-73922: The Proxy server settings are not working correctly for the Telemetry Publisher in Cloudera Manager versions 7.11.3 and higher.

The Proxy server issues are resolved by updating the cdp-sdk-java artifact's version. This issue is now resolved.

OPSAPS-73792: Telemetry Publisher exhibited incorrect behaviour during job uploads by accepting a Status Code 503 response and marking logs as successfully exported.

The issue is now resolved. Telemetry Publisher now treats only Status Code 200 as successful. For non-200 status codes, Telemetry Publisher will now log an error message.

OPSAPS-73655: Cloud replication no longer fails after the delegation token is issued

You can now configure com.cloudera.enterprise.distcp.skip-delegation-token-on-cloud-replication = false in the Cloudera Manager Clusters *HDFS SERVICE* Configuration HDFS Replication

Advanced Configuration Snippet (Safety Valve) for core-site.xml property to ensure that the HDFS and Hive external table replication policies replicating from an on-premises cluster to cloud do not fail.

The replication policies were failing when you chose the Advanced Setting Delete Policy Delete permanently option during the replication policy creation process and an incremental replication run was in progress.

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-73142: The required configuration from replication safety valve is not accessed

An Ozone replication policy with Incremental with fallback to full file listing option failed with Pre-Filelisting Check Failed with Error: target bucket has layout OBS, but [fs.s3a.endpoint, fs.s3a.secret.key, fs.s3a.access.key] properties are missing from the target Ozone service core-site.xml config error because the required configuration was not available in the required folders.

To mitigate this issue, the required configuration parameters are now added automatically to the required folders during the Ozone replication policy run.

OPSAPS-73219, OPSAPS-73218: Dry run of Ozone incremental policies fail

When you run the Cloudera Manager API request to start an Ozone replication policy in Dry Run mode, the replication policy fails if the OzoneReplicationType is Incremental only or Incremental with fallback to full file listing. To prevent this issue, the Dry Run operation is no longer available.

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

Some remote replication commands continue to run endlessly even after a Cloudera Manager restart operation. This issue is fixed.

OPSAPS-73158, OPSAPS-73902: HDFS replication policies fail when the policies prefetch the expired Kerberos ticket from the 'sourceTicketCache' file

To ensure that the replication policies do not prefetch the expired Kerberos ticket from the sourceTicketCache file before the replication policy run, you must add the 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.

Fixed Common Vulnerabilities and Exposures

For information about Common Vulnerabilities and Exposures (CVE) that are fixed in Cloudera Manager 7.11.3 cumulative hotfix 16, see Fixed Common Vulnerabilities and Exposures in Cloudera Manager 7.11.3 cumulative hotfixes.

The repositories for Cloudera Manager 7.11.3 CHF16 are listed in the following table:

Table 2: Cloudera Manager 7.11.3 CHF16

Repository Type	Repository Location

RHEL 9 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.36/redhat9/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.36/redhat9/yum/cloudera-manager.repo`

Repository Type	Repository Location

RHEL 8 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.36/redhat8/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.36/redhat8/yum/cloudera-manager.repo`

RHEL 7 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.36/redhat7/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.36/redhat7/yum/cloudera-manager.repo`

SLES 15	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.36/sles15/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.36/sles15/yum/cloudera-manager.repo`

SLES 12	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.36/sles12/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.36/sles12/yum/cloudera-manager.repo`

Ubuntu 22	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.36/ubuntu2204/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.36/ubuntu2204/apt/cloudera-manager.list`

Repository Type	Repository Location

Ubuntu 20	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.36/ubuntu2004/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.36/ubuntu2004/apt/cloudera-manager.list`

Cloudera Manager 7.11.3 Cumulative hotfix 15

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 15.

This cumulative hotfix was released on May 8, 2025.

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

New features and changed behavior for Cloudera Manager 7.11.3 CHF 15 (version: 7.11.3.34-66004006):
OPSAPS-71124: Swapping full Perl stack with Perl interpreter package

The Cloudera Manager Agent now depends on the Perl interpreter package rather than the full Perl stack. As a result, the GCC toolchain is not pulled in, so no compiler components are installed with the Cloudera Manager Agent. There is no functional impact to Cloudera Manager Agents.

OPSAPS-70909: 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.

OPSAPS-73164: Ozone's upgrade handlers are not properly added to the UpgradeHandlerRegistry

Certain upgrade handlers are not added anymore during an upgrade, but this change in behaviour corrects potential problems by skipping correctly the UpgradeHandlers that are not designated to run for a certain upgrade path.

OPSAPS-73075: Add 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.

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF 15 (version: 7.11.3.34-66004006):
OPSAPS-75899: HDFS directory creation fails on JDK 11 or higher when LDAP or Active Directory integrated clusters using the hadoop.security.group.mapping property.

Due to this issue, many critical Cloudera Manager operations fail to complete, such as:

- Install Oozie ShareLib (Oozie Actions Install Oozie ShareLib)
- Install YARN MapReduce Framework JARs (YARN Actions Install YARN MapReduce Framework JARs)

You must perform the following workaround steps to manually add specific Java modules to the HDFS service script on all nodes in the cluster:

1. Navigate to the directory /opt/cloudera/cm-agent/service/hdfs/.
2. Open the hdfs.sh file for editing.
3. Locate the line starting with MKDIR_JAVA_OPTS.
4. Update it to include the following flags:

```
Change this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=f
ile://${MKDIR_LOG4J_FILE}"
To this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=fi
le://${MKDIR_LOG4J_FILE} --add-exports=java.naming/com.sun.j
ndi.ldap=ALL-UNNAMED --add-opens=java.naming/com.sun.jndi.ld
ap=ALL-UNNAMED"
```

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

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.

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

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

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-73211: Cloudera Manager 7.11.3 does not clean up Python Path impacting Hue to start

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.11.3 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-73655: Cloud replication fails after the delegation token is issued

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-72984: Alerts due to change in hostname fetching functionality in jdk 8 and jdk 11

Upgrading JAVA from JDK 8 to JDK 11 creates the following alert in CMS:

Bad : CMSERVER:pit666.slayer.mayank: Reaching Cloudera Manager Server failed

This happens due to a functionality change in JDK 11 on hostname fetching.

```
[root@pit666.slayer ~]# /us/lib/jvm/java-1.8.0/bin/java GetHostN
ame
```

```
Hostname: pit666.slayer.mayank

[root@pit666.slayer ~]# /usr/lib/jvm/java-11/bin/java GetHostName
Hostname: pit666.slayer
```

You can notice that the "hostname" is set to a short name instead of FQDN.

The current workaround is to set the hostname as FQDN.

OPSAPS-72784: Upgrades from CDH6 to CDP Private Cloud Base 7.1.9 SP1 or higher versions fail with a health check timeout exception

If you are using Cloudera Manager 7.11.3 cumulative hotfix 14 or higher versions and upgrading from CDH 6 to CDP Private Cloud Base 7.1.9 SP1 or higher versions, the upgrade fails with a CMUpgradeHealthException timeout exception. This is because upgrades from CDH 6 to CDP Private Cloud Base 7.1.9 SP1 or to any of its cumulative hotfix versions are not supported.

None.

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

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

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

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

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

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

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

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

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

OPSAPS-70771: Running replication policy runs must not allow you to download the performance reports

During a replication policy run, the A server error has occurred. See Cloudera Manager server log for details error message appears on the UI and the Cloudera Manager log shows "java.lang.IllegalStateException: Command has no result data." when you click:

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

This is because the Replication Manager UI shows the performance report links as enabled and clickable which is incorrect. You can download the reports only after the replication job run is complete.

None

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

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

None

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

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

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-74398: Ozone and HDFS replication policies might fail when you use different destination proxy user and source proxy user

HDFS on-premises to on-premises replication fails when the following conditions are true:

- You configure different Run As Username and Run on Peer as Username during the replication policy creation process.
- The user configured in Run As Username does not have the permission to access the source path on the source HDFS.

Ozone replication fails when the following conditions are true:

- FSO-to-FSO replication or an OBS-to-OBS replication with Incremental with fallback to full file listing or Incremental only replication type.
- You configured different Run As Username and Run on Peer as Username during the replication policy creation process.
- The user configured in Run As Username does not have the permission to access the source bucket on the source Ozone.

Provide the same permissions to the user configured in Run As Username as the permissions of Run on Peer as Username on the source cluster.

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

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.

Following are the list of fixed issues that were shipped for Cloudera Manager 7.11.3 CHF 15 (version: 7.11.3.34-66004006):

OPSAPS-73164: Ozone's upgrade handlers are not properly added to the UpgradeHandlerRegistry

Ozone Upgrade handlers were not properly applied in certain CDP upgrade scenarios. This fix corrects potential problems by skipping correctly the UpgradeHandlers that are not designated to run for a certain upgrade path.

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 expo rt CMF_SERVER_ARGS=), so the parameter cannot be utilized correctly.

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-69383: HTTP header used the wrong Strict_Transport_Security header

Previously, the HADOOP_HTTP_HEADER_STRICT_TRANSPORT_SECURITY parameter used the wrong header:hadoop.http.header.Strict_Transport_Security syntax. This issue is now resolved and the HTTP header name is now corrected to Strict-Transport-Security.

OPSAPS-70983: Hive replication command fails for Sentry to Ranger replication

Hive replication command for Sentry to Ranger replication works as expected now. 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-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-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-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-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-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.

OPSAPS-73481: Knox readiness check gateway-status endpoint should return the list of topologies for which it is waiting for

Knox readiness check for gateway-status endpoint now returns the list of topologies for which it is waiting for. Before the update you had to check the gateway.log to understand what are topologies Knox is waiting for to be deployed.

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

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.

Fixed Common Vulnerabilities and Exposures

For information about Common Vulnerabilities and Exposures (CVE) that are fixed in Cloudera Manager 7.11.3 cumulative hotfix 15, see Fixed Common Vulnerabilities and Exposures in Cloudera Manager 7.11.3 cumulative hotfixes.

The repositories for Cloudera Manager 7.11.3-CHF 15 are listed in the following table:

Table 3: Cloudera Manager 7.11.3-CHF 15

Repository Type	Repository Location

RHEL 9 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.34/redhat9/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.34/redhat9/yum/cloudera-manager.repo`

RHEL 8 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.34/redhat8/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.34/redhat8/yum/cloudera-manager.repo`

RHEL 7 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.34/redhat7/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.34/redhat7/yum/cloudera-manager.repo`

Repository Type	Repository Location

SLES 15	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.34/sles15/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.34/sles15/yum/cloudera-manager.repo`

SLES 12	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.34/sles12/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.34/sles12/yum/cloudera-manager.repo`

Ubuntu 22	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.34/ubuntu2204/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.34/ubuntu2204/apt/cloudera-manager.list`

Ubuntu 20	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.34/ubuntu2004/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.34/ubuntu2004/apt/cloudera-manager.list`

Cloudera Manager 7.11.3 Cumulative hotfix 14

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 14.

This cumulative hotfix was released on April 7, 2025.

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

New features and changed behavior for Cloudera Manager 7.11.3 CHF 14 (version: 7.11.3.33-64908687): OPSAPS-73151: Improve Ranger Admin Diagnostic Collection command from Cloudera Manager scripts

The Ranger Admin Diagnostic Collection command is enhanced and a new configuration option called `ranger.admin.diag.metrics.collection.type` is introduced. This option allows you to specify the type of metrics data to be collected.

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF 14 (version: 7.11.3.33-64908687):

OPSAPS-75899: HDFS directory creation fails on JDK 11 or higher when LDAP or Active Directory integrated clusters using the hadoop.security.group.mapping property.

Due to this issue, many critical Cloudera Manager operations fail to complete, such as:

- Install Oozie ShareLib (Oozie Actions Install Oozie ShareLib)
- Install YARN MapReduce Framework JARs (YARN Actions Install YARN MapReduce Framework JARs)

You must perform the following workaround steps to manually add specific Java modules to the HDFS service script on all nodes in the cluster:

1. Navigate to the directory /opt/cloudera/cm-agent/service/hdfs/.
2. Open the hdfs.sh file for editing.
3. Locate the line starting with MKDIR_JAVA_OPTS.
4. Update it to include the following flags:

```
Change this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=f
ile://${MKDIR_LOG4J_FILE}"
To this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=fi
le://${MKDIR_LOG4J_FILE} --add-exports=java.naming/com.sun.j
ndi.ldap=ALL-UNNAMED --add-opens=java.naming/com.sun.jndi.ld
ap=ALL-UNNAMED"
```

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

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.

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

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

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-73211: Cloudera Manager 7.11.3 does not clean up Python Path impacting Hue to start

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.11.3 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-72984: Alerts due to change in hostname fetching functionality in jdk 8 and jdk 11

Upgrading JAVA from JDK 8 to JDK 11 creates the following alert in CMS:

Bad : CMSERVER:pit666.slayer.mayank: Reaching Cloudera Manager Server failed

This happens due to a functionality change in JDK 11 on hostname fetching.

```
[root@pit666.slayer ~]# /us/lib/jvm/java-1.8.0/bin/java GetHostN
ame
Hostname: pit666.slayer.mayank
```

```
[root@pit666.slayer ~]# /usr/lib/jvm/java-11/bin/java GetHostName
Hostname: pit666.slayer
```

You can notice that the "hostname" is set to a short name instead of FQDN.

The current workaround is to set the hostname as FQDN.

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.

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 resume the replication policies' job run.
- Upgrade source Cloudera Manager, upgrade target Cloudera Manager, and rerun the failed Ozone replication policies between the source and target clusters.

OPSAPS-72784: Upgrades from CDH6 to CDP Private Cloud Base 7.1.9 SP1 or higher versions fail with a health check timeout exception

If you are using Cloudera Manager 7.11.3 cumulative hotfix 14 and upgrading from CDH 6 to CDP Private Cloud Base 7.1.9 SP1 or higher versions, the upgrade fails with a CMUpgradeHealthExcep tion timeout exception. This is because upgrades from CDH 6 to CDP Private Cloud Base 7.1.9 SP1 or to any of its cumulative hotfix versions are not supported.

None.

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 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-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 7.11.3 or Cloudera Manager 7.11.3 CHF-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.

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

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

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

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

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

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

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

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

OPSAPS-70771: Running replication policy runs must not allow you to download the performance reports

During a replication policy run, the A server error has occurred. See Cloudera Manager server log for details error message appears on the UI and the Cloudera Manager log shows "java.lang.IllegalStateException: Command has no result data." when you click:

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

This is because the Replication Manager UI shows the performance report links as enabled and clickable which is incorrect. You can download the reports only after the replication job run is complete.

None

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

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

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

None

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

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

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-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.

None

OPSAPS-74398: Ozone and HDFS replication policies might fail when you use different destination proxy user and source proxy user

HDFS on-premises to on-premises replication fails when the following conditions are true:

- You configure different Run As Username and Run on Peer as Username during the replication policy creation process.
- The user configured in Run As Username does not have the permission to access the source path on the source HDFS.

Ozone replication fails when the following conditions are true:

- FSO-to-FSO replication or an OBS-to-OBS replication with Incremental with fallback to full file listing or Incremental only replication type.
- You configured different Run As Username and Run on Peer as Username during the replication policy creation process.
- The user configured in Run As Username does not have the permission to access the source bucket on the source Ozone.

Provide the same permissions to the user configured in Run As Username as the permissions of Run on Peer as Username on the source cluster.

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

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.

Following are the list of fixed issues that were shipped for Cloudera Manager 7.11.3 CHF 14 (version: 7.11.3.33-64908687):

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

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. This issue is fixed.

OPSAPS-71635: NoSuchElementException appears if Hive tables are deleted during Hive export process

During a Hive external table replication policy run, if the Hive tables were deleted during the Hive export step, the "NoSuchElementException" appeared and the replication policy run failed.

This issue is fixed. Replication Manager now shows this exception as a TABLE_NOTFOUND_ERROR, and ignores the deleted tables during the replication policy run.

CDPD-79831, HIVE-27797: The timed out transactions are not logged as 'ABORTED' in NOTIFICATION_LOG

The timed out transactions were not getting logged or marked as aborted in the notification_log table. This issue is fixed.

CDPD-53185, HIVE-28772: REPL_TXN_MAP table on target cluster is not cleared a Hive ACID replication policy is deleted

The entry in the 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 specified in the replication policy get deleted on the target cluster even before the transaction is completely replicated.

In this scenario, if you drop the policy and database and perform a re-bootstrap on the same database then the replicated database on the target cluster is tagged as 'database incompatible' after hive.repl.txn.timeout (default is 11 days). This happens after the housekeeper thread process deletes the retained entry. This issue is fixed.

Fixed Common Vulnerabilities and Exposures

For information about Common Vulnerabilities and Exposures (CVE) that are fixed in Cloudera Manager 7.11.3 cumulative hotfix 14, see Fixed Common Vulnerabilities and Exposures in Cloudera Manager 7.11.3 cumulative hotfixes.

The repositories for Cloudera Manager 7.11.3-CHF 14 are listed in the following table:

Table 4: Cloudera Manager 7.11.3-CHF 14

Repository Type	Repository Location

RHEL 9 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.33/redhat9/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.33/redhat9/yum/cloudera-manager.repo`

RHEL 8 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.33/redhat8/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.33/redhat8/yum/cloudera-manager.repo`

Repository Type	Repository Location

RHEL 7 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.33/redhat7/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.33/redhat7/yum/cloudera-manager.repo`

SLES 15	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.33/sles15/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.33/sles15/yum/cloudera-manager.repo`

SLES 12	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.33/sles12/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.33/sles12/yum/cloudera-manager.repo`

Ubuntu 22	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.33/ubuntu2204/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.33/ubuntu2204/apt/cloudera-manager.list`

Ubuntu 20	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.33/ubuntu2004/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.33/ubuntu2004/apt/cloudera-manager.list`

Cloudera Manager 7.11.3 Cumulative hotfix 13

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 13.

This cumulative hotfix was released on March 7, 2025.

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

New features and changed behavior for Cloudera Manager 7.11.3 CHF 13 (version: 7.11.3.32-63522208):
OPSAPS-69339: Deleting VERSION file, bootstrap file, certificates and keys after OM decommissioning

After running the Ozone Manager decommissioning command, the VERSION file, bootstrap file, certificates and keys are deleted.

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF 13 (version: 7.11.3.32-63522208):
OPSAPS-75899: HDFS directory creation fails on JDK 11 or higher when LDAP or Active Directory integrated clusters using the hadoop.security.group.mapping property.

Due to this issue, many critical Cloudera Manager operations fail to complete, such as:

- Install Oozie ShareLib (Oozie Actions Install Oozie ShareLib)
- Install YARN MapReduce Framework JARs (YARN Actions Install YARN MapReduce Framework JARs)

You must perform the following workaround steps to manually add specific Java modules to the HDFS service script on all nodes in the cluster:

1. Navigate to the directory /opt/cloudera/cm-agent/service/hdfs/.
2. Open the hdfs.sh file for editing.
3. Locate the line starting with MKDIR_JAVA_OPTS.
4. Update it to include the following flags:

```
Change this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=f
ile://${MKDIR_LOG4J_FILE}"
To this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=fi
le://${MKDIR_LOG4J_FILE} --add-exports=java.naming/com.sun.j
ndi.ldap=ALL-UNNAMED --add-opens=java.naming/com.sun.jndi.ld
ap=ALL-UNNAMED"
```

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

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.

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

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

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-73211: Cloudera Manager 7.11.3 does not clean up Python Path impacting Hue to start

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.11.3 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-72984: Alerts due to change in hostname fetching functionality in jdk 8 and jdk 11

Upgrading JAVA from JDK 8 to JDK 11 creates the following alert in CMS:

Bad : CMSERVER:pit666.slayer.mayank: Reaching Cloudera Manager Server failed

This happens due to a functionality change in JDK 11 on hostname fetching.

```
[root@pit666.slayer ~]# /us/lib/jvm/java-1.8.0/bin/java GetHostN
ame
Hostname: pit666.slayer.mayank

[root@pit666.slayer ~]# /usr/lib/jvm/java-11/bin/java GetHostName
Hostname: pit666.slayer
```

You can notice that the "hostname" is set to a short name instead of FQDN.

The current workaround is to set the hostname as FQDN.

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 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-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 7.11.3 or Cloudera Manager 7.11.3 CHF-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-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.

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 resume the replication policies' job runs.
- Upgrade source Cloudera Manager, upgrade target Cloudera Manager, and rerun the failed Ozone replication policies between the source and target clusters.

OPSAPS-72784: Upgrades from CDH6 to CDP Private Cloud Base 7.1.9 SP1 or higher versions fail with a health check timeout exception

If you are using Cloudera Manager 7.11.3 cumulative hotfix 13 and upgrading from CDH 6 to CDP Private Cloud Base 7.1.9 SP1 or higher versions, the upgrade fails with a CMUpgradeHealthException timeout exception. This is because upgrades from CDH 6 to CDP Private Cloud Base 7.1.9 SP1 or to any of its cumulative hotfix versions are not supported.

None.

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

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

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

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

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

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

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

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

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

OPSAPS-70771: Running replication policy runs must not allow you to download the performance reports

During a replication policy run, the A server error has occurred. See Cloudera Manager server log for details error message appears on the UI and the Cloudera Manager log shows "java.lang.IllegalStateException: Command has no result data." when you click:

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

This is because the Replication Manager UI shows the performance report links as enabled and clickable which is incorrect. You can download the reports only after the replication job run is complete.

None

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

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

None

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

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

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

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

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

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

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

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

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

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-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.

None

OPSAPS-74398: Ozone and HDFS replication policies might fail when you use different destination proxy user and source proxy user

HDFS on-premises to on-premises replication fails when the following conditions are true:

- You configure different Run As Username and Run on Peer as Username during the replication policy creation process.
- The user configured in Run As Username does not have the permission to access the source path on the source HDFS.

Ozone replication fails when the following conditions are true:

- FSO-to-FSO replication or an OBS-to-OBS replication with Incremental with fallback to full file listing or Incremental only replication type.
- You configured different Run As Username and Run on Peer as Username during the replication policy creation process.
- The user configured in Run As Username does not have the permission to access the source bucket on the source Ozone.

Provide the same permissions to the user configured in Run As Username as the permissions of Run on Peer as Username on the source cluster.

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

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-75090: Ozone replication policies fail without source proxy user

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.

Following are the list of fixed issues that were shipped for Cloudera Manager 7.11.3 CHF 13 (version: 7.11.3.32-63522208):

OPSAPS-71527: Hive Metrics not loading after installing Cloudera Manager Cumulative hotfix - 7.11.3.9 (CHF6)

After applying a Cloudera Manager Cumulative hotfix - 7.11.3.9 (CHF6) upgrade, multiple Hive metric charts showed no data, specifically for data collected post the Cloudera Manager upgrade timestamp.

A user would see this problem when attempting an upgrade from a lower Cumulative hotfix (for example, CM 7.11.3 CHF1 (7.11.3.2) to a higher Cumulative hotfix within the same release line (for example, CM 7.11.3 CHF6 (CM 7.11.3.9).

This issue is fixed in CM 7.11.3 CHF13 (7.11.3.32) and higher versions. This fix resolves an issue where Hive metric charts failed to display data after applying a Cloudera Manager Cumulative hotfix upgrade.

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-72632: Cloudera Manager - Stale service restart API call is failing

When there is a configuration change for the Cloudera Management Service (CMS), process staleness detection for the CMS does not work. This issue is fixed now.

OPSAPS-71933: Telemetry Publisher is unable to publish Spark event logs to 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-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-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.

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

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. 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 Manager 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.

Fixed Common Vulnerabilities and Exposures

For information about Common Vulnerabilities and Exposures (CVE) that are fixed in Cloudera Manager 7.11.3 cumulative hotfix 13, see Fixed Common Vulnerabilities and Exposures in Cloudera Manager 7.11.3 cumulative hotfixes.

The repositories for Cloudera Manager 7.11.3-CHF 13 are listed in the following table:

Table 5: Cloudera Manager 7.11.3-CHF 13

Repository Type	Repository Location

RHEL 9 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.32/redhat9/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.32/redhat9/yum/cloudera-manager.repo`

Repository Type	Repository Location

RHEL 8 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.32/redhat8/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.32/redhat8/yum/cloudera-manager.repo`

RHEL 7 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.32/redhat7/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.32/redhat7/yum/cloudera-manager.repo`

SLES 15	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.32/sles15/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.32/sles15/yum/cloudera-manager.repo`

SLES 12	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.32/sles12/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.32/sles12/yum/cloudera-manager.repo`

Ubuntu 22	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.32/ubuntu2204/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.32/ubuntu2204/apt/cloudera-manager.list`

Repository Type	Repository Location

Ubuntu 20	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.32/ubuntu2004/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.32/ubuntu2004/apt/cloudera-manager.list`

Cloudera Manager 7.11.3 Cumulative hotfix 12

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 12.

This cumulative hotfix was released on February 20, 2025.

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

New features and changed behavior for Cloudera Manager 7.11.3 CHF 12 (version: 7.11.3.31-62995507):
OPSAPS-71872: FedRAMP-Compliant TLS Cipher Configuration for Kudu

Earlier, the default TLS ciphers for Kudu were not FedRAMP compliant and could not be configured through Cloudera Manager. To address this issue, the default TLS cipher values have been updated to align with FedRAMP compliance, and Kudu now allows configuring the minimum TLS version and cipher suite preferences directly through Cloudera Manager. This enhancement ensures improved security and compliance for TLS-secured RPC connections.

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF 12 (version: 7.11.3.31-62995507):
OPSAPS-75899: HDFS directory creation fails on JDK 11 or higher when LDAP or Active Directory integrated clusters using the hadoop.security.group.mapping property.

Due to this issue, many critical Cloudera Manager operations fail to complete, such as:

- Install Oozie ShareLib (Oozie Actions Install Oozie ShareLib)
- Install YARN MapReduce Framework JARs (YARN Actions Install YARN MapReduce Framework JARs)

You must perform the following workaround steps to manually add specific Java modules to the HDFS service script on all nodes in the cluster:

1. Navigate to the directory /opt/cloudera/cm-agent/service/hdfs/.
2. Open the hdfs.sh file for editing.
3. Locate the line starting with MKDIR_JAVA_OPTS.
4. Update it to include the following flags:

```
Change this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=f
ile://${MKDIR_LOG4J_FILE}"
To this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=fi
le://${MKDIR_LOG4J_FILE} --add-exports=java.naming/com.sun.j
ndi.ldap=ALL-UNNAMED --add-opens=java.naming/com.sun.jndi.ld
ap=ALL-UNNAMED"
```

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

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.

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

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

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-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.

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 resume the replication policies' job runs.
- Upgrade source Cloudera Manager, upgrade target Cloudera Manager, and rerun the failed Ozone replication policies between the source and target clusters.

OPSAPS-73211: Cloudera Manager 7.11.3 does not clean up Python Path impacting Hue to start

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.11.3 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-73655: Cloud replication fails after the delegation token is issued

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-72984: Alerts due to change in hostname fetching functionality in jdk 8 and jdk 11

Upgrading JAVA from JDK 8 to JDK 11 creates the following alert in CMS:

Bad : CMSERVER:pit666.slayer.mayank: Reaching Cloudera Manager Server failed

This happens due to a functionality change in JDK 11 on hostname fetching.

```
[root@pit666.slayer ~]# /us/lib/jvm/java-1.8.0/bin/java GetHostN
ame
```

```
Hostname: pit666.slayer.mayank

[root@pit666.slayer ~]# /usr/lib/jvm/java-11/bin/java GetHostName
Hostname: pit666.slayer
```

You can notice that the "hostname" is set to a short name instead of FQDN.

The current workaround is to set the hostname as FQDN.

OPSAPS-72784: Upgrades from CDH6 to CDP Private Cloud Base 7.1.9 SP1 or higher versions fail with a health check timeout exception

If you are using Cloudera Manager 7.11.3 cumulative hotfix 12 and upgrading from CDH 6 to CDP Private Cloud Base 7.1.9 SP1 or higher versions, the upgrade fails with a CMUpgradeHealthExcep tion timeout exception. This is because upgrades from CDH 6 to CDP Private Cloud Base 7.1.9 SP1 or to any of its cumulative hotfix versions are not supported.

None.

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 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-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 7.11.3 or Cloudera Manager 7.11.3 CHF-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.

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

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

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

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

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

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

MariaDB 10.6, by default, includes the property require_secure_transport=ON in the configuration file (/etc/my.cnf), which is absent in MariaDB 10.4. This setting prohibits non-TLS connections,

leading to access issues. This problem is observed in High Availability (HA) mode, where certain operations may not be using the same connection.

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

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

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

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

None

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

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

None

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

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

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

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

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

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

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

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-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.

None

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

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

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

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

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-74398: Ozone and HDFS replication policies might fail when you use different destination proxy user and source proxy user

HDFS on-premises to on-premises replication fails when the following conditions are true:

- You configure different Run As Username and Run on Peer as Username during the replication policy creation process.
- The user configured in Run As Username does not have the permission to access the source path on the source HDFS.

Ozone replication fails when the following conditions are true:

- FSO-to-FSO replication or an OBS-to-OBS replication with Incremental with fallback to full file listing or Incremental only replication type.
- You configured different Run As Username and Run on Peer as Username during the replication policy creation process.

- The user configured in Run As Username does not have the permission to access the source bucket on the source Ozone.

Provide the same permissions to the user configured in Run As Username as the permissions of Run on Peer as Username on the source 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.

None

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

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.

Following are the list of fixed issues that were shipped for Cloudera Manager 7.11.3 CHF 12 (version: 7.11.3.31-62995507):

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-72215: ECS Cloudera Manager 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.

When you want to update docker cert through Cloudera Manager UI config, add '\\n' to specify a newline character in the certificate.

For example:

```
-----BEGIN CERTIFICATE-----\nMIIERTCCAy2gAwIBAgIUIL8o1MjD5he7nZ
KKa/C8rx9uPjcwDQYJKoZIhvcNAQEL\nBQAwXTELMAkGA1
UEBhMCVVMxEzARBgNVBAgMCkNhbGlmb3JuaWExEzARBgNVBAcM\nClNhbnRhQ2xhc
mExETAPBgN
VBAoMCENsb3VkZXJhMREwDwYDVQQLDAhDbG91ZGVy\nYTAeFw0yNDAzMTExMjU5ND
VaFw0zNDA
zMDkxMjU5NDVaMF0xCzAJBgNVBAYTAlVT\nMRMwEQYDVQQIDApDYWxpZm9ybml
hMRMwEQYDVQ
QHDApTYW50YUNsYXJhMREwDwYD\nVQQKDAhDbG91ZGVyYTERMA8GA1UECwwIQ2xv
dWRlcmEwg
gEiMA0GCSqGSIb3DQEB\nAQUAA4IBDwAwggEKAoIBAQDcuxGszWmzVnWCwDICn
lxUBtO+tI6RPs2jx
Q7C7kIj\nTHTaQ2kGl/ZzQOJBpYT/jFmiQGPSKb4iLSxed+Xk5xAOkNWDIL+Hlf
5txjkw/FTf\nHiyWep9Da
QDF07M/Cl3nb8JmpRyA5fKYpVbJAFIEXOhTxrcnH/4o5ubLM7mHVXwY\nafoPD
5AuiOD/I+xxmqb/x+fKt
HzY1eEzDb2vjjDJBRqxpHvg/S4hHsgZJ7wU7wg+\nPk4uPV3O83h9NI+b4SOwXu
nuKRCCh4dRKm8/Q
w4f7tDFdCAIubvO1AGtfyJJp9xR\npMIjhIuna1K2TnPQomdoIy/KqrFFzVaHevy
inEnRLG2NAgMBAAGjgfw
```

```
wgfkwHQYD\nVR0OBBYEFHWX21/BhL5J5kNpxmb8FmDchlmBMIGaBgNVHSMEgZIwgY
+AFHWX21/B
\nhL5J5kNpxmb8FmDchlmBoWGkXzBdMQswCQYDVQQGEwJVUzETMBEGA1UECAwKQ
2Fs\naWZvcm
5pYTETMBEGA1UEBwwKU2FudGFDbGFyYTERMA8GA1UECgwIQ2xvdWRlcmEx\nETAP
BgNVBAsMCEN
sb3VkZXJhghQgvyjUyMPmF7udkopr8LyvH24+NzAMBgNVHRME\nBTADAQH/MAsGA
1UdDwQEAwIC/
DAPBgNVHREECDAGhwQKgW26MA8GA1UdEgQIMAaH\nBAqBbbowDQYJKoZIhvcNAQ
ELBQADggEBA
Mks+sY+ETaPzFLg2PolUTT4GeXEqnGl\nSmZzIkiA6l2DCYQD/7mTLd5Ea63oI78
fxatRnG5MLf5aHVLs4
W+WYhoP6B7HLPUo\nNGPJviRBHtUDRYVpD5Q0hhQtHB4Q1H+sgrE53VmbIQqLPOAx
vpM//oJCFDT8
NbOI\n+bTJ48N34ujosjNaiP6x09xbzRzOnYd6VyhZ/pgsiRZ4qlZsVyv1TImP9
VpHcC7P\nukxNuBdXBS3j
EXcvEV1Eq4Di+z6PIWoPIHUunQ9P0akYEvbXuL88knM5FNhS6YBP\nGd91KkGd
z6srRIVRiF+XP0e6IwZC70kkWiw
f8vX/CuR64ZQxc30ot70=\n-----END CERTIFICATE-----\n
```

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.

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-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-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 does not load completely. This issue is fixed.

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-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-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.

OPSAPS-72111: Directory creation fails during Hive ACID replication policy creation if the target cluster uses Dell EMC Isilon storage

Directory creation failed during the Hive ACID replication policy creation process if the target cluster used Dell EMC Isilon storage. To mitigate this issue, ensure that the *hive* user and the *hive* group have 0755 port permission to the staging location.

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-71105: Expose or set YARN cgroup v2 settings in Cloudera Manager

Cgroup v2 support is now enabled by default, and YARN detects and uses the correct cgroup handling code.

OPSAPS-72427: Node Managers fail to start with the No cgroup controllers file found error

Previously, when cgroup was enabled, cgroup v2 was enabled automatically. This caused RM startup failures, in case of cgroup v1-only clusters, due to a missing cgroup.controllers file. This issue is now resolved and cgroup v2 support now falls back to v1 when there are no v2 controllers.

Note: This issue was fixed in runtime and was merged to 7.1.9 SP1 CHF 5. You must update CDP runtime to CDP 7.1.9 SP1 CHF 5 along with the Cloudera Manager update to CM 7.11.3 CHF12 or higher.

Fixed Common Vulnerabilities and Exposures

For information about Common Vulnerabilities and Exposures (CVE) that are fixed in Cloudera Manager 7.11.3 cumulative hotfix 12, see Fixed Common Vulnerabilities and Exposures in Cloudera Manager 7.11.3 cumulative hotfixes.

The repositories for Cloudera Manager 7.11.3-CHF 12 are listed in the following table:

Table 6: Cloudera Manager 7.11.3-CHF 12

Repository Type	Repository Location

RHEL 9 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.31/redhat9/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.31/redhat9/yum/cloudera-manager.repo`

RHEL 8 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.31/redhat8/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.31/redhat8/yum/cloudera-manager.repo`

RHEL 7 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.31/redhat7/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.31/redhat7/yum/cloudera-manager.repo`

SLES 15	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.31/sles15/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.31/sles15/yum/cloudera-manager.repo`

Repository Type	Repository Location

SLES 12	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.31/sles12/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.31/sles12/yum/cloudera-manager.repo`

Ubuntu 22	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.31/ubuntu2204/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.31/ubuntu2204/apt/cloudera-manager.list`

Ubuntu 20	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.31/ubuntu2004/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.31/ubuntu2004/apt/cloudera-manager.list`

Cloudera Manager 7.11.3 Cumulative hotfix 11

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 11.

This cumulative hotfix was released on December 19, 2024.

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

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF 11 (version: 7.11.3.28-60766845):
OPSAPS-75899: HDFS directory creation fails on JDK 11 or higher when LDAP or Active Directory integrated clusters using the hadoop.security.group.mapping property.

Due to this issue, many critical Cloudera Manager operations fail to complete, such as:

- Install Oozie ShareLib (Oozie Actions Install Oozie ShareLib)
- Install YARN MapReduce Framework JARs (YARN Actions Install YARN MapReduce Framework JARs)

You must perform the following workaround steps to manually add specific Java modules to the HDFS service script on all nodes in the cluster:

1. Navigate to the directory /opt/cloudera/cm-agent/service/hdfs/.
2. Open the hdfs.sh file for editing.
3. Locate the line starting with MKDIR_JAVA_OPTS.

4. Update it to include the following flags:

```
Change this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=f
ile://${MKDIR_LOG4J_FILE}"
To this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=fi
le://${MKDIR_LOG4J_FILE} --add-exports=java.naming/com.sun.j
ndi.ldap=ALL-UNNAMED --add-opens=java.naming/com.sun.jndi.ld
ap=ALL-UNNAMED"
```

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

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.

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

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

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-73211: Cloudera Manager 7.11.3 does not clean up Python Path impacting Hue to start

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.11.3 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

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-72984: Alerts due to change in hostname fetching functionality in jdk 8 and jdk 11

Upgrading JAVA from JDK 8 to JDK 11 creates the following alert in CMS:

Bad : CMSERVER:pit666.slayer.mayank: Reaching Cloudera Manager Server failed

This happens due to a functionality change in JDK 11 on hostname fetching.

```
[root@pit666.slayer ~]# /us/lib/jvm/java-1.8.0/bin/java GetHostN
ame
Hostname: pit666.slayer.mayank

[root@pit666.slayer ~]# /usr/lib/jvm/java-11/bin/java GetHostName
Hostname: pit666.slayer
```

You can notice that the "hostname" is set to a short name instead of FQDN.

The current workaround is to set the hostname as FQDN.

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 7.11.3 or Cloudera Manager 7.11.3 CHF-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.

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

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

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

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

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

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

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

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

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

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

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

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

None

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

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

None

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

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

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

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

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

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

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

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

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

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-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.

None

OPSAPS-72509, CDPD-32440: Hive metadata transfer to GCS fails with ClassNotFoundException

Hive replication policies from an on-premises cluster to cloud fails during the “Transfer Metadata Files” step if the following conditions are true:

- the target is a GCS Data Lake
- 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 is because the fs.gs.delegation.token.binding property is already defined in the configuration and cannot be unset to disable the delegation tokens in the cloud connector service.

None

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

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

None

OPSAPS-72427: Node Managers fail to start for cgroup v1-only clusters

Enabling cgroup automatically enables cgroup v2. In case of cgroup v1-only clusters, there is RM startup failures due to a missing cgroup.controllers file.

To temporarily resolve this issue on cgroup v1 clusters, the v2 support must be turned off by adding the yarn.nodemanager.linux-container-executor.cgroups.v2.enabled property with the value false to the NodeManager Advanced Configuration Snippet (Safety Valve) for yarn-site.xml.

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.

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 resume the replication policies' job run.
- Upgrade source Cloudera Manager, upgrade target Cloudera Manager, and rerun the failed Ozone replication policies between the source and target clusters.

OPSAPS-74398: Ozone and HDFS replication policies might fail when you use different destination proxy user and source proxy user

HDFS on-premises to on-premises replication fails when the following conditions are true:

- You configure different Run As Username and Run on Peer as Username during the replication policy creation process.
- The user configured in Run As Username does not have the permission to access the source path on the source HDFS.

Ozone replication fails when the following conditions are true:

- FSO-to-FSO replication or an OBS-to-OBS replication with Incremental with fallback to full file listing or Incremental only replication type.
- You configured different Run As Username and Run on Peer as Username during the replication policy creation process.

- The user configured in Run As Username does not have the permission to access the source bucket on the source Ozone.

Provide the same permissions to the user configured in Run As Username as the permissions of Run on Peer as Username on the source cluster.

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

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.

Following are the list of fixed issues that were shipped for Cloudera Manager 7.11.3 CHF 11 (version: 7.11.3.28-60766845):

OPSAPS-66996: When Cloudera Manager is creating home directories for service accounts, it will skip creating directories for those accounts that are already defined through an external LDAP or SAML service connected to the operating system on the host. This occurs because the attempt to create the service user account fails, causing follow on actions such as home directory creation to be skipped.

Cloudera Manager now detects when the external authentication services connected to the host operating system (through PAM or SSSD or other methods) have pre-created service accounts and will proceed with service account home directory creation if the directory is missing on the local host filesystem.

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

This issue is now resolved by changing the default configuration value on Cloudera Manager. Now the ozone.snapshot.deep.cleaning.enabled is disabled by default and ozone.snapshot.ordered.deletion.enabled enabled by default on Cloudera Manager.

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

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 now resolved.

OPSAPS-72105: Kafka, SRM, and SMM cannot process messages compressed with Zstd or Snappy if /tmp is mounted as noexec

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

OPSAPS-71931: Ranger-HDFS plugin resource lookup is not working with JDK 17 on Isilon cluster

The issue is fixed now by adding the sun.net.util package to Ranger Admin Java opts for jdk 17.

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.

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, and the required validation on the source completes successfully as expected.

OPSAPS-70752: The "MapReduce Service" field shows incorrect services during the Ranger replication creation process

This issue is fixed. The MapReduce service field shows the MapReduce services in the destination cluster when you choose the Replicate Ranger audit logs in HDFS option in the **Create Ranger replication policy** wizard in Replication Manager.

OPSAPS-72229: Atlas replication policies consider active and inactive Atlas server instances during the replication policy runs

This issue is resolved. Replication Manager now only considers the active Atlas server instances during the Atlas replication policy runs.

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 resolved.

OPSAPS-72208: Set YARN cgroup v2 settings in Cloudera Manager

Cgroup v2 support is now enabled by default. This support is backward compatible, therefore cgroup v1 is automatically detected and used if required.

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 now resolved and when there are no applications running in a queue, then its capacity is set to zero.

Fixed Common Vulnerabilities and Exposures

For information about Common Vulnerabilities and Exposures (CVE) that are fixed in Cloudera Manager 7.11.3 cumulative hotfix 11, see Fixed Common Vulnerabilities and Exposures in Cloudera Manager 7.11.3 cumulative hotfixes.

The repositories for Cloudera Manager 7.11.3-CHF 11 are listed in the following table:

Table 7: Cloudera Manager 7.11.3-CHF 11

Repository Type	Repository Location

RHEL 9 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.28/redhat9/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.28/redhat9/yum/cloudera-manager.repo`

Repository Type	Repository Location

RHEL 8 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.28/redhat8/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.28/redhat8/yum/cloudera-manager.repo`

RHEL 7 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.28/redhat7/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.28/redhat7/yum/cloudera-manager.repo`

SLES 15	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.28/sles15/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.28/sles15/yum/cloudera-manager.repo`

SLES 12	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.28/sles12/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.28/sles12/yum/cloudera-manager.repo`

Ubuntu 20	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.28/ubuntu2004/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.28/ubuntu2004/apt/cloudera-manager.list`

Repository Type	Repository Location

Ubuntu 22	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.28/ubuntu2204/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.28/ubuntu2204/apt/cloudera-manager.list`

Cloudera Manager 7.11.3 Cumulative hotfix 10

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 10.

This cumulative hotfix was released on November 12, 2024.

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

New features and changed behavior for Cloudera Manager 7.11.3 CHF 10 (version: 7.11.3.26-58725444): Enhancements to the Observability page

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

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

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF 10 (version: 7.11.3.26-58725444):
OPSAPS-75899: HDFS directory creation fails on JDK 11 or higher when LDAP or Active Directory integrated clusters using the hadoop.security.group.mapping property.

Due to this issue, many critical Cloudera Manager operations fail to complete, such as:

- Install Oozie ShareLib (Oozie Actions Install Oozie ShareLib)
- Install YARN MapReduce Framework JARs (YARN Actions Install YARN MapReduce Framework JARs)

You must perform the following workaround steps to manually add specific Java modules to the HDFS service script on all nodes in the cluster:

1. Navigate to the directory /opt/cloudera/cm-agent/service/hdfs/.
2. Open the hdfs.sh file for editing.
3. Locate the line starting with MKDIR_JAVA_OPTS.
4. Update it to include the following flags:

```
Change this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=f
ile://${MKDIR_LOG4J_FILE}"
To this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=fi
le://${MKDIR_LOG4J_FILE} --add-exports=java.naming/com.sun.j
ndi.ldap=ALL-UNNAMED --add-opens=java.naming/com.sun.jndi.ld
ap=ALL-UNNAMED"
```

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

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.

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

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

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-73211: Cloudera Manager 7.11.3 does not clean up Python Path impacting Hue to start

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.11.3 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

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-72984: Alerts due to change in hostname fetching functionality in jdk 8 and jdk 11

Upgrading JAVA from JDK 8 to JDK 11 creates the following alert in CMS:

Bad : CMSERVER:pit666.slayer.mayank: Reaching Cloudera Manager Server failed

This happens due to a functionality change in JDK 11 on hostname fetching.

```
[root@pit666.slayer ~]# /us/lib/jvm/java-1.8.0/bin/java GetHostN
ame
Hostname: pit666.slayer.mayank

[root@pit666.slayer ~]# /usr/lib/jvm/java-11/bin/java GetHostName
Hostname: pit666.slayer
```

You can notice that the "hostname" is set to a short name instead of FQDN.

The current workaround is to set the hostname as FQDN.

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

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-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 7.11.3 or Cloudera Manager 7.11.3 CHF-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.

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

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

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

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

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

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

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

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

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

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

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

None

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

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

None

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

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

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

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

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

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

During the Ozone replication policy run, Replication Manager does not read the value in the 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-72468: Subsequent Ozone OBS-to-OBS replication policy do not skip replicated files during replication

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

None

OPSAPS-72509, CDPD-32440: Hive metadata transfer to GCS fails with ClassNotFoundException

Hive replication policies from an on-premises cluster to cloud fails during the "Transfer Metadata Files" step if the following conditions are true:

- the target is a GCS Data Lake
- 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 is because the fs.gs.delegation.token.binding property is already defined in the configuration and cannot be unset to disable the delegation tokens in the cloud connector service.

None

OPSAPS-74398: Ozone and HDFS replication policies might fail when you use different destination proxy user and source proxy user

HDFS on-premises to on-premises replication fails when the following conditions are true:

- You configure different Run As Username and Run on Peer as Username during the replication policy creation process.
- The user configured in Run As Username does not have the permission to access the source path on the source HDFS.

Ozone replication fails when the following conditions are true:

- FSO-to-FSO replication or an OBS-to-OBS replication with Incremental with fallback to full file listing or Incremental only replication type.

- You configured different Run As Username and Run on Peer as Username during the replication policy creation process.
- The user configured in Run As Username does not have the permission to access the source bucket on the source Ozone.

Provide the same permissions to the user configured in Run As Username as the permissions of Run on Peer as Username on the source cluster.

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

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.

Following are the list of fixed issues that were shipped for Cloudera Manager 7.11.3 CHF 10 (version: 7.11.3.26-58725444):

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-70704: Kerberos connectivity check does not work as expected with JDK17 when you add Cloudera Manager peers

When you add a source Cloudera Manager that supports JDK17, the Kerberos connectivity check fails and the Error while reading /etc/krb5.conf on <hostname ; for all hosts>... error message appears. This issue is fixed now.

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-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-71615: Service monitor crashes due to an out-of-memory (OOM) error.

This issue is fixed now.

OPSAPS-71258: Kafka, SRM, and SMM cannot process messages compressed with Zstd or Snappy

Kafka, Streams Replication Manager, and Streams Messaging Manager can now process messages compressed with Zstd and Snappy if /tmp is mounted with the noexec option.

This fix changes the default Zstd and Snappy temporary directory from /tmp to the following service specific directories.

- Kafka - /var/lib/kafka
- Streams Messaging Manager - /var/lib/streams_messaging_manager
- Streams Replication Manager - /var/lib/streams_messaging_manager

Ensure that each service user has write and execute permission on the directory specific to their service. Otherwise, the service will fail to process compressed messages.

- The Kafka service user (default: kafka) must have write and execute permission on /var/lib/kafka.
- The Streams Messaging Manager service user (default: streamsmsgmgr) must have write and execute permission on /var/lib/streams_messaging_manager.
- The Streams Replication Manager service user (default: streamsrepmgr) must have write and execute permission on /var/lib/streams_replication_manager.

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 resolved.

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-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-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 runs these policies concurrently.

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

CDPD-85189: Open transactions from target cluster after Hive ACID table incremental replication job creates issues

After every incremental replication job run for the Hive ACID table replication policy, the transactions that remain open after the dump and load operation now get aborted on the target cluster.

This is applicable to only those transactions that were not in an open state during the dump operation on the source cluster.

CDPD-84058: Commit message does not provide the required details

Previously, the commit event message did not provide details about whether the event was associated with the read or write operation. The commit text in the notification_log message is now enhanced, so that you can use the commit details during the dump process to decide whether you want to skip or include the commit during the next replication job run.

Fixed Common Vulnerabilities and Exposures

For information about Common Vulnerabilities and Exposures (CVE) that are fixed in Cloudera Manager 7.11.3 cumulative hotfix 10, see Fixed Common Vulnerabilities and Exposures in Cloudera Manager 7.11.3 cumulative hotfixes.

The repositories for Cloudera Manager 7.11.3-CHF 10 are listed in the following table:

Table 8: Cloudera Manager 7.11.3-CHF 10

Repository Type	Repository Location

RHEL 9 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.26/redhat9/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.26/redhat9/yum/cloudera-manager.repo`

RHEL 8 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.26/redhat8/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.26/redhat8/yum/cloudera-manager.repo`

RHEL 7 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.26/redhat7/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.26/redhat7/yum/cloudera-manager.repo`

SLES 15	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.26/sles15/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.26/sles15/yum/cloudera-manager.repo`

Repository Type	Repository Location

SLES 12	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.26/sles12/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.26/sles12/yum/cloudera-manager.repo`

Ubuntu 20	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.26/ubuntu2004/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.26/ubuntu2004/apt/cloudera-manager.list`

Ubuntu 22	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.26/ubuntu2204/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.26/ubuntu2204/apt/cloudera-manager.list`

Cloudera Manager 7.11.3 Cumulative hotfix 9.1

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 9.1.

This cumulative hotfix was released on October 10, 2024.

Important: Cloudera Manager 7.11.3 CHF 9.1 (7.11.3.24) has replaced Cloudera Manager CHF 9 (7.11.3.21). Contact Cloudera Support for questions related to any specific hotfixes.

New features and changed behavior for Cloudera Manager 7.11.3 CHF 9.1 (version: 7.11.3.24-58365749): Enhancements to the Observability page

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

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

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF 9.1 (version: 7.11.3.24-58365749):
OPSAPS-75899: HDFS directory creation fails on JDK 11 or higher when LDAP or Active Directory integrated clusters using the hadoop.security.group.mapping property.

Due to this issue, many critical Cloudera Manager operations fail to complete, such as:

- Install Oozie ShareLib (Oozie Actions Install Oozie ShareLib)

- Install YARN MapReduce Framework JARs (YARN Actions Install YARN MapReduce Framework JARs)

You must perform the following workaround steps to manually add specific Java modules to the HDFS service script on all nodes in the cluster:

1. Navigate to the directory /opt/cloudera/cm-agent/service/hdfs/.
2. Open the hdfs.sh file for editing.
3. Locate the line starting with MKDIR_JAVA_OPTS.
4. Update it to include the following flags:

```
Change this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=f
ile://${MKDIR_LOG4J_FILE}"
To this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=fi
le://${MKDIR_LOG4J_FILE} --add-exports=java.naming/com.sun.j
ndi.ldap=ALL-UNNAMED --add-opens=java.naming/com.sun.jndi.ld
ap=ALL-UNNAMED"
```

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

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.

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

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

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-73211: Cloudera Manager 7.11.3 does not clean up Python Path impacting Hue to start

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.11.3 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

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-72984: Alerts due to change in hostname fetching functionality in jdk 8 and jdk 11

Upgrading JAVA from JDK 8 to JDK 11 creates the following alert in CMS:

Bad : CMSERVER:pit666.slayer.mayank: Reaching Cloudera Manager Server failed

This happens due to a functionality change in JDK 11 on hostname fetching.

```
[root@pit666.slayer ~]# /us/lib/jvm/java-1.8.0/bin/java GetHostN
ame
Hostname: pit666.slayer.mayank

[root@pit666.slayer ~]# /usr/lib/jvm/java-11/bin/java GetHostName
Hostname: pit666.slayer
```

You can notice that the "hostname" is set to a short name instead of FQDN.

The current workaround is to set the hostname as FQDN.

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 7.11.3 or Cloudera Manager 7.11.3 CHF-x versions. When you run the Host Inspector, you get an error Not finding Psycopg2, even though it is installed on all hosts.

None

TSB 2024-806: Important upgrade required for Cloudera Manager versions 7.11.3 CHF8 and 7.11.3 CHF9

Cloudera recommends that customers running versions 7.11.3 with Cumulative Hotfix (CHF) 9 or 7.11.3 CHF8 of Cloudera Manager should upgrade to version 7.11.3 CHF9.1 to prevent issues with cluster management.

For the latest update on this issue see the corresponding Knowledge article: TSB 2024-806: Important upgrade required for Cloudera Manager versions 7.11.3 CHF8 and 7.11.3 CHF9.

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.

Manually add = sign to file_metadata_reload_properties configuration and modify the Gflags configuration file when the file_metadata_reload_properties configuration is passed as empty.

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

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

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

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

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

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

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

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

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

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

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

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

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

None

OPSAPS-70704: Kerberos connectivity check does not work as expected with JDK17 when you add Cloudera Manager peers

When you add a source Cloudera Manager that supports JDK17, the Kerberos connectivity check fails and the Error while reading /etc/krb5.conf on <hostname ; for all hosts>... error message appears.

None

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

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

None

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

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

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

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

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

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

During the Ozone replication policy run, Replication Manager does not read the value in the 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-72509, CDPD-32440: Hive metadata transfer to GCS fails with ClassNotFoundException

Hive replication policies from an on-premises cluster to cloud fails during the "Transfer Metadata Files" step if the following conditions are true:

- the target is a GCS Data Lake
- 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 is because the fs.gs.delegation.token.binding property is already defined in the configuration and cannot be unset to disable the delegation tokens in the cloud connector service.

None

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

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.

Following are the list of fixed issues that were shipped for Cloudera Manager 7.11.3 CHF 9.1 (version: 7.11.3.24-58365749):

OPSAPS-70915: Cloudera Manager Agent incorrectly detected its own cgroup path

The Cloudera Manager Agent incorrectly detected its own cgroup path and created the YARN NodeManager's cgroup (for example, hadoop-yarn) under the Cloudera Manager Agent's system.slice hierarchy instead of the root-level cgroup path.

For example, the YARN cgroup appeared as /sys/fs/cgroup/cpu,cpuacct/system.slice/cloudera-scm-agent.service/hadoop-yarn instead of /sys/fs/cgroup/cpu,cpuacct/hadoop-yarn.

Because of this misplacement, when you restart the Cloudera Manager Agent process, systemd automatically destroys the nested cgroup (/system.slice/cloudera-scm-agent.service/hadoop-yarn). This immediately kills all running YARN containers and causes active jobs to fail. This issue is fixed now.

With this fix, the Cloudera Manager Agent now:

- Automatically detects and uses the correct (root-level) cgroup paths.
- No longer depends on systemd's service slice hierarchy to determine cgroup locations.
- Ensures that restarting the Cloudera Manager Agent does not impact running YARN containers or jobs.

However, if you previously added manual workarounds (for example, uncommenting or hardcoding cgroup paths in /etc/cloudera-scm-agent/config.ini), you can safely retain them—but you no longer need them after this fix.

OPSAPS-71249: Auto Action trigger for Impala Engine fails

Auto action triggers for the Impala engine do not work for Kerberos-enabled Private Cloud Base clusters. This issue is fixed now.

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-71210: Ozone Basic Canary displays an exception about loading S3 secret from keystore java.lang.RuntimeException: Encountered error when loading S3 secret from keystore: java.lang.NullPointerException.

This issue is now resolved.

OPSAPS-69603: Ozone CLI is not available to the CMON role if CMON is not installed on the same cluster as CDH.. This results in failure of Ozone Basic Canary because canary uses Ozone CLI to access Ozone.

This issue is now resolved. Ozone Basic Canary now runs successfully even if CMON and Ozone roles/gateway are running on different hosts or clusters.

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

The schedule frequency works as expected after you edit the existing Ozone replication policies.

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-71615: Service monitor crashes due to an out-of-memory (OOM) error.

This issue is now resolved.

Fixed Common Vulnerabilities and Exposures

For information about Common Vulnerabilities and Exposures (CVE) that are fixed in Cloudera Manager 7.11.3 cumulative hotfix 9.1, see Fixed Common Vulnerabilities and Exposures in Cloudera Manager 7.11.3 cumulative hotfixes.

The repositories for Cloudera Manager 7.11.3-CHF 9.1 are listed in the following table:

Table 9: Cloudera Manager 7.11.3-CHF 9.1

Repository Type	Repository Location

RHEL 9 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.24/redhat9/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.24/redhat9/yum/cloudera-manager.repo`

RHEL 8 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.24/redhat8/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.24/redhat8/yum/cloudera-manager.repo`

RHEL 7 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.24/redhat7/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.24/redhat7/yum/cloudera-manager.repo`

SLES 15	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.24/sles15/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.24/sles15/yum/cloudera-manager.repo`

SLES 12	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.24/sles12/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.24/sles12/yum/cloudera-manager.repo`

Repository Type	Repository Location

Ubuntu 20	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.24/ubuntu2004/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.24/ubuntu2004/apt/cloudera-manager.list`

Ubuntu 22	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.24/ubuntu2204/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.24/ubuntu2204/apt/cloudera-manager.list`

Cloudera Manager 7.11.3 Cumulative hotfix 8

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 8.

This cumulative hotfix was released on August 27, 2024.

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

New features and changed behavior for Cloudera Manager 7.11.3 CHF8 (version: 7.11.3.16-56304673):
Added ability in the Cloudera Manager Agent's config.ini file to disable filesystem checks.

In Cloudera Manager Agent 7.11.3 CHF8 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.

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF8 (version: 7.11.3.16-56304673):
OPSAPS-75899: HDFS directory creation fails on JDK 11 or higher when LDAP or Active Directory integrated clusters using the hadoop.security.group.mapping property.

Due to this issue, many critical Cloudera Manager operations fail to complete, such as:

- Install Oozie ShareLib (Oozie Actions Install Oozie ShareLib)
- Install YARN MapReduce Framework JARs (YARN Actions Install YARN MapReduce Framework JARs)

You must perform the following workaround steps to manually add specific Java modules to the HDFS service script on all nodes in the cluster:

1. Navigate to the directory /opt/cloudera/cm-agent/service/hdfs/.
2. Open the hdfs.sh file for editing.
3. Locate the line starting with MKDIR_JAVA_OPTS.
4. Update it to include the following flags:

```
Change this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=f
ile://${MKDIR_LOG4J_FILE}"
To this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=fi
le://${MKDIR_LOG4J_FILE} --add-exports=java.naming/com.sun.j
ndi.ldap=ALL-UNNAMED --add-opens=java.naming/com.sun.jndi.ld
ap=ALL-UNNAMED"
```

OPSAPS-70915: Cloudera Manager Agent incorrectly detected its own cgroup path

The Cloudera Manager Agent incorrectly detected its own cgroup path and created the YARN NodeManager's cgroup (for example, hadoop-yarn) under the Cloudera Manager Agent's system.slice hierarchy instead of the root-level cgroup path.

For example, the YARN cgroup appeared as /sys/fs/cgroup/cpu,cpuacct/system.slice/cloudera-scm-agent.service/hadoop-yarn instead of /sys/fs/cgroup/cpu,cpuacct/hadoop-yarn.

Because of this misplacement, when you restart the Cloudera Manager Agent process, systemd automatically destroys the nested cgroup (/system.slice/cloudera-scm-agent.service/hadoop-yarn). This immediately kills all running YARN containers and causes active jobs to fail.

To prevent YARN from inheriting the Cloudera Manager Agent's cgroup hierarchy, you can explicitly configure Cloudera Manager Agent to use the root cgroup path. Perform this configuration by uncommenting and setting the cgroups paths in the Cloudera Manager Agent configuration file: /etc/cloudera-scm-agent/config.ini. Perform the following steps:

1. Under the [cgroups] section, uncomment or add the following lines (adjusting for your cgroup version and controller types):

```
[cgroups]
mounts=cpu,cpuacct,cpuset,memory
cpu_cgroup_mount_point=/sys/fs/cgroup/cpu,cpuacct
memory_cgroup_mount_point=/sys/fs/cgroup/memory
```

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

```
sudo systemctl restart cloudera-scm-agent
```

This configuration ensures that the Cloudera Manager Agent and its managed roles (such as YARN NodeManager) always use root-level cgroup paths rather than inheriting them from system.slice, and prevents systemd from automatically cleaning up those cgroups when Cloudera Manager Agent restarts.

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

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.

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

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

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-73211: Cloudera Manager 7.11.3 does not clean up Python Path impacting Hue to start

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.11.3 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

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-72984: Alerts due to change in hostname fetching functionality in jdk 8 and jdk 11

Upgrading JAVA from JDK 8 to JDK 11 creates the following alert in CMS:

Bad : CMSERVER:pit666.slayer.mayank: Reaching Cloudera Manager Server failed

This happens due to a functionality change in JDK 11 on hostname fetching.

```
[root@pit666.slayer ~]# /us/lib/jvm/java-1.8.0/bin/java GetHostN
ame
Hostname: pit666.slayer.mayank

[root@pit666.slayer ~]# /usr/lib/jvm/java-11/bin/java GetHostName
Hostname: pit666.slayer
```

You can notice that the "hostname" is set to a short name instead of FQDN.

The current workaround is to set the hostname as FQDN.

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 7.11.3 or Cloudera Manager 7.11.3 CHF-x versions. When you run the Host Inspector, you get an error Not finding Psycopg2, even though it is installed on all hosts.

None

TSB 2024-806: Important upgrade required for Cloudera Manager versions 7.11.3 CHF8 and 7.11.3 CHF9

Cloudera recommends that customers running versions 7.11.3 with Cumulative Hotfix (CHF) 9 or 7.11.3 CHF8 of Cloudera Manager should upgrade to version 7.11.3 CHF9.1 to prevent issues with cluster management.

For the latest update on this issue see the corresponding Knowledge article: TSB 2024-806: Important upgrade required for Cloudera Manager versions 7.11.3 CHF8 and 7.11.3 CHF9.

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.

Manually add = sign to file_metadata_reload_properties configuration and modify the Gflags configuration file when the file_metadata_reload_properties configuration is passed as empty.

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

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

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

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

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

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

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

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

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

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

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

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

None

OPSAPS-70704: Kerberos connectivity check does not work as expected with JDK17 when you add Cloudera Manager peers

When you add a source Cloudera Manager that supports JDK17, the Kerberos connectivity check fails and the Error while reading /etc/krb5.conf on <hostname ; for all hosts>... error message appears.

None

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

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

None

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

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

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

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

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

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

During the Ozone replication policy run, Replication Manager does not read the value in the 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-72509, CDPD-32440: Hive metadata transfer to GCS fails with ClassNotFoundException

Hive replication policies from an on-premises cluster to cloud fails during the "Transfer Metadata Files" step if the following conditions are true:

- the target is a GCS Data Lake
- 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 is because the fs.gs.delegation.token.binding property is already defined in the configuration and cannot be unset to disable the delegation tokens in the cloud connector service.

None

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

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

Following are the list of fixed issues that were shipped for Cloudera Manager 7.11.3 CHF8 (version: 7.11.3.16-53216725):

OPSAPS-68845: Cloudera Manager Server fails to start after the Cloudera Manager upgrade

Starting from the Cloudera Manager 7.11.3 version up to the Cloudera Manager 7.11.3 CHF7 version, the Cloudera Manager Server fails to start after the Cloudera Manager upgrade due to Navigator user roles improperly handled in the upgrade in some scenarios. This issue is fixed now by removing the extra Navigator roles.

OPSAPS-70419: The Livy3 server lacks necessary iceberg configurations in spark-defaults.

Attempting to query the Iceberg table using Livy3 failed with Error in loading storage han dler.org.apache.iceberg.mr.hive.HiveIcebergStorageHandler. The same query was successful when executed using spark3-shell.

Now Iceberg is added to the Livy3 classpath.

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-70831: Regenerate missing keytab option listed for ECS/Docker instances, even though it is no-op

On the Cloudera Manager UI, the Regenerate missing Keytab option is now hidden for Docker and ECS instances and this option remains visible for all other service types.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

OPSAPS-68353: Ozone Canary in Cloudera Manager Service Monitor uses keystore to store S3 secret.

Ozone Basic Canary now uses the S3 secret stored in the Cloudera Manager keystore instead of sending a request to Ozone. If the S3 secret is not available in the keystore, a request is sent to Ozone for the S3 secret credentials and is stored in the keystore.

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-71005: RemoteCmdWork is using a singlethreaded executor

By default, Replication Manager runs the remote commands for a replication policy through a single-thread executor. You can search and enable the enable_multithreaded_remote_cmd_executor property in the target Cloudera Manager Administration Settings page to run future replication policies through the multi-threaded executor. This action improves the processing performance of the replication workloads.

Additionally, you can also change the multithreaded_remote_cmd_executor_max_threads and multithreaded_remote_cmd_executor_keepalive_time properties to fine-tune the replication policy performance.

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.

Fixed Common Vulnerabilities and Exposures

For information about Common Vulnerabilities and Exposures (CVE) that are fixed in Cloudera Manager 7.11.3 cumulative hotfix 8, see Fixed Common Vulnerabilities and Exposures in Cloudera Manager 7.11.3 cumulative hotfixes.

The repositories for Cloudera Manager 7.11.3-CHF8 are listed in the following table:

Table 10: Cloudera Manager 7.11.3-CHF8

Repository Type	Repository Location

RHEL 9 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.16/redhat9/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.16/redhat9/yum/cloudera-manager.repo`

RHEL 8 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.16/redhat8/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.16/redhat8/yum/cloudera-manager.repo`

Repository Type	Repository Location

RHEL 7 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.16/redhat7/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.16/redhat7/yum/cloudera-manager.repo`

SLES 15	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.16/sles15/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.16/sles15/yum/cloudera-manager.repo`

SLES 12	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.16/sles12/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.16/sles12/yum/cloudera-manager.repo`

Ubuntu 20	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.16/ubuntu2004/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.16/ubuntu2004/apt/cloudera-manager.list`

Ubuntu 22	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.16/ubuntu2204/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.16/ubuntu2204/apt/cloudera-manager.list`

Cloudera Manager 7.11.3 Cumulative hotfix 7

Know more about the Cloudera Manager 7.11.3 cumulative hotfix 7.

Version for CDP Private Cloud Base customers: 7.11.3.11.

Version for CDP Private Cloud Data Services customers: 7.11.3.14.

This cumulative hotfix was released on July 19, 2024.

Important: Cloudera Manager 7.11.3 CHF7 and later versions support CDP Private Cloud Base 7.1.9 SP1.

Important: Cloudera Manager 7.11.3 CHF7 Data Services (version: 7.11.3.14) supports CDP Private Cloud Data Services 1.5.4 CHF1 release.

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

New features

For information about new and changed or updated features, see What's New in Cloudera Manager 7.11.3 Cumulative hotfix 7 (CDP Private Cloud Base 7.1.9 SP1).

New platform support

For information about new platform support enhancements, see What's new in Platform Support.

Known issues

For information about the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF7, see Known Issues in Cloudera Manager 7.11.3 Cumulative hotfix 7 (CDP Private Cloud Base 7.1.9 SP1).

Fixed issues

For information about the list of fixed issues that are shipped for Cloudera Manager 7.11.3 CHF7, see Fixed Issues in Cloudera Manager 7.11.3 Cumulative hotfix 7 (CDP Private Cloud Base 7.1.9 SP1).

Fixed Common Vulnerabilities and Exposures

For information about Common Vulnerabilities and Exposures (CVE) that are fixed in Cloudera Manager 7.11.3 cumulative hotfix 7, see Fixed Common Vulnerabilities and Exposures in Cloudera Manager 7.11.3 cumulative hotfix 7.

Deprecation notices

For information about the deprecated or removed operating systems and databases from the Cloudera Manager 7.11.3 Cumulative hotfix 7, see Platform and OS.

The repositories for Cloudera Manager 7.11.3-CHF7 are listed in the following table:

Table 11: Cloudera Manager 7.11.3-CHF7 (Version: 7.11.3.14)

Repository Type	Repository Location

RHEL 9 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.14/redhat9/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.14/redhat9/yum/cloudera-manager.repo`

Repository Type	Repository Location

RHEL 8 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.14/redhat8/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.14/redhat8/yum/cloudera-manager.repo`

RHEL 7 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.14/redhat7/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.14/redhat7/yum/cloudera-manager.repo`

Table 12: Cloudera Manager 7.11.3-CHF7 (Version: 7.11.3.11)

Repository Type	Repository Location

RHEL 9 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.11/redhat9/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.11/redhat9/yum/cloudera-manager.repo`

RHEL 8 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.11/redhat8/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.11/redhat8/yum/cloudera-manager.repo`

RHEL 7 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.11/redhat7/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.11/redhat7/yum/cloudera-manager.repo`

Repository Type	Repository Location

SLES 15	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.11/sles15/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.11/sles15/yum/cloudera-manager.repo`

SLES 12	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.11/sles12/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.11/sles12/yum/cloudera-manager.repo`

Ubuntu 20	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.11/ubuntu2004/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.11/ubuntu2004/apt/cloudera-manager.list`

Ubuntu 22	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.11/ubuntu2204/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.11/ubuntu2204/apt/cloudera-manager.list`

Cloudera Manager 7.11.3 Cumulative hotfix 6

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 6.

This cumulative hotfix was released on May 28, 2024.

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

New features and changed behavior for Cloudera Manager 7.11.3 CHF6 (version: 7.11.3.9-53216725):

ECS Update Ingress Controller Certificate action is now available through the API

You can now access the ECS Update Ingress Controller Certificate action through the API and the UI.

Password protection support for Ingress private key

Ingress certificate private key is now supported with password protection.

Deploy client configuration command timed-out on larger node clusters

The Deploy Client Config command is improved now. Previously, it could take long and time-out on large clusters. It is now leveraging multithreading and optimized for parallel execution. The command is now expected to complete much faster and should not cause timeouts.

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF6 (version: 7.11.3.9-53216725):
OPSAPS-75899: HDFS directory creation fails on JDK 11 or higher when LDAP or Active Directory integrated clusters using the hadoop.security.group.mapping property.

Due to this issue, many critical Cloudera Manager operations fail to complete, such as:

- Install Oozie ShareLib (Oozie Actions Install Oozie ShareLib)
- Install YARN MapReduce Framework JARs (YARN Actions Install YARN MapReduce Framework JARs)

You must perform the following workaround steps to manually add specific Java modules to the HDFS service script on all nodes in the cluster:

1. Navigate to the directory /opt/cloudera/cm-agent/service/hdfs/.
2. Open the hdfs.sh file for editing.
3. Locate the line starting with MKDIR_JAVA_OPTS.
4. Update it to include the following flags:

```
Change this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=f
ile://${MKDIR_LOG4J_FILE}"
To this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=fi
le://${MKDIR_LOG4J_FILE} --add-exports=java.naming/com.sun.j
ndi.ldap=ALL-UNNAMED --add-opens=java.naming/com.sun.jndi.ld
ap=ALL-UNNAMED"
```

OPSAPS-70915: Cloudera Manager Agent incorrectly detected its own cgroup path

The Cloudera Manager Agent incorrectly detected its own cgroup path and created the YARN NodeManager's cgroup (for example, hadoop-yarn) under the Cloudera Manager Agent's system.slice hierarchy instead of the root-level cgroup path.

For example, the YARN cgroup appeared as /sys/fs/cgroup/cpu,cpuacct/system.slice/cloudera-scm-agent.service/hadoop-yarn instead of /sys/fs/cgroup/cpu,cpuacct/hadoop-yarn.

Because of this misplacement, when you restart the Cloudera Manager Agent process, systemd automatically destroys the nested cgroup (/system.slice/cloudera-scm-agent.service/hadoop-yarn). This immediately kills all running YARN containers and causes active jobs to fail.

To prevent YARN from inheriting the Cloudera Manager Agent's cgroup hierarchy, you can explicitly configure Cloudera Manager Agent to use the root cgroup path. Perform this configuration by uncommenting and setting the cgroups paths in the Cloudera Manager Agent configuration file: /etc/cloudera-scm-agent/config.ini. Perform the following steps:

1. Under the [cgroups] section, uncomment or add the following lines (adjusting for your cgroup version and controller types):

```
[cgroups]
mounts=cpu,cpuacct,cpuset,memory
cpu_cgroup_mount_point=/sys/fs/cgroup/cpu,cpuacct
memory_cgroup_mount_point=/sys/fs/cgroup/memory
```

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

```
sudo systemctl restart cloudera-scm-agent
```

This configuration ensures that the Cloudera Manager Agent and its managed roles (such as YARN NodeManager) always use root-level cgroup paths rather than inheriting them from syst em.slice, and prevents systemd from automatically cleaning up those cgroups when Cloudera Manager Agent restarts.

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

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.

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

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

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-73211: Cloudera Manager 7.11.3 does not clean up Python Path impacting Hue to start

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.11.3 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

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-72984: Alerts due to change in hostname fetching functionality in jdk 8 and jdk 11

Upgrading JAVA from JDK 8 to JDK 11 creates the following alert in CMS:

Bad : CMSERVER:pit666.slayer.mayank: Reaching Cloudera Manager Server failed

This happens due to a functionality change in JDK 11 on hostname fetching.

```
[root@pit666.slayer ~]# /us/lib/jvm/java-1.8.0/bin/java GetHostN
ame
Hostname: pit666.slayer.mayank

[root@pit666.slayer ~]# /usr/lib/jvm/java-11/bin/java GetHostName
Hostname: pit666.slayer
```

You can notice that the "hostname" is set to a short name instead of FQDN.

The current workaround is to set the hostname as FQDN.

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 7.11.3 or Cloudera Manager 7.11.3 CHF-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-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.

Manually add = sign to file_metadata_reload_properties configuration and modify the Gflags configuration file when the file_metadata_reload_properties configuration is passed as empty.

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.

To resolve this issue, disable the Enable Log Event Capture configuration on the Configuration page of a service.

OPSAPS-68845: Cloudera Manager Server fails to start after the Cloudera Manager upgrade

Starting from the Cloudera Manager 7.11.3 version up to the Cloudera Manager 7.11.3 CHF7 version, the Cloudera Manager Server fails to start after the Cloudera Manager upgrade due to Navigator user roles improperly handled in the upgrade in some scenarios.

None

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.

Upgrade to CDP 7.1.9, or downgrade to Cloudera Manager 7.7.1.

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

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

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

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

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

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

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

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

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

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

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

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

Azul Open JDK 8

RHEL or SLES

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

Ubuntu or Debian

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

Azul Open JDK 11

For DEBs only

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

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

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

Upgrade the source Cloudera Manager to 7.11.3 or higher version.

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

The status of the deleted table displays as ACTIVE.

None

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

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

None

Following are the list of fixed issues that were shipped for Cloudera Manager 7.11.3 CHF6 (version: 7.11.3.9-53216725):

OPSAPS-70207: Cloudera Manager Agents sending the Impala profile data with an incorrect header

Fixed an issue where Impala query profile data was not visible in Cloudera Observability with Cloudera Manager Agents running on Python 3. This was caused by Cloudera Manager Agents sending the profile data with an incorrect Content-Type in the HTTP header. Telemetry Publisher responded to the issue with incorrect Content-Type stopping the data flow. The issue was fixed by correcting the Content-Type header. No further action is required.

OPSAPS-65460: The current RetryWrapper implementation does not work as expected when the transient database error appears

Hive replication policies no longer fail with the javax.persistence.OptimisticLockException error on the source cluster during the Hive export step.

OPSAPS-60832: Decommissioning process for HDFS DataNodes is not completed in Cloudera Manager

This issue has been fixed by renewing the Kerberos ticket, which addresses Kerberos expiration issues during the DataNode decommissioning process.

OPSAPS-68418: Partition missing during column statistics import operation

A data-race issue found during the Hive metadata export step during the Hive external table replication policy run has been fixed so that concurrent modifications made to the partitions during the export operation does not result in import failure.

OPSAPS-70079: NPE appears during the directory creation process in the Isilon clusters because the sourceRoleForKerberos value is null

After you configure the Dell EMC Isilon clusters, you must ensure that you configure the following options on the Cloudera Manager Administration Settings page:

- Custom Kerberos Keytab Location (to be used for replication for secure clusters on this Cloudera Manager)
- Custom Kerberos Principal Name (to be used for replication for secure clusters on this Cloudera Manager)

Tip: Use the configured Custom Kerberos Principal Name value in the Run As Username field during the replication policy creation process when using Isilon storage clusters. For more information, see How to resolve replication policies that fail with the “Custom keytab configuration is required for this service” error.

The repositories for Cloudera Manager 7.11.3-CHF6 are listed in the following table:

Table 13: Cloudera Manager 7.11.3-CHF6

Repository Type	Repository Location

RHEL 9 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.9/redhat9/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.9/redhat9/yum/cloudera-manager.repo`

RHEL 8 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.9/redhat8/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.9/redhat8/yum/cloudera-manager.repo`

Repository Type	Repository Location

RHEL 7 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.9/redhat7/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.9/redhat7/yum/cloudera-manager.repo`

SLES 15	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.9/sles15/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.9/sles15/yum/cloudera-manager.repo`

SLES 12	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.9/sles12/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.9/sles12/yum/cloudera-manager.repo`

Ubuntu 20	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.9/ubuntu2004/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.9/ubuntu2004/apt/cloudera-manager.list`

Cloudera Manager 7.11.3 Cumulative hotfix 5

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 5.

This cumulative hotfix was released on April 8, 2024.

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

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF5 (version: 7.11.3.7-52024171):
OPSAPS-75899: HDFS directory creation fails on JDK 11 or higher when LDAP or Active Directory integrated clusters using the hadoop.security.group.mapping property.

Due to this issue, many critical Cloudera Manager operations fail to complete, such as:

- Install Oozie ShareLib (Oozie Actions Install Oozie ShareLib)
- Install YARN MapReduce Framework JARs (YARN Actions Install YARN MapReduce Framework JARs)

You must perform the following workaround steps to manually add specific Java modules to the HDFS service script on all nodes in the cluster:

1. Navigate to the directory /opt/cloudera/cm-agent/service/hdfs/.
2. Open the hdfs.sh file for editing.
3. Locate the line starting with MKDIR_JAVA_OPTS.
4. Update it to include the following flags:

```
Change this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=f
ile://${MKDIR_LOG4J_FILE}"
To this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=fi
le://${MKDIR_LOG4J_FILE} --add-exports=java.naming/com.sun.j
ndi.ldap=ALL-UNNAMED --add-opens=java.naming/com.sun.jndi.ld
ap=ALL-UNNAMED"
```

OPSAPS-70915: Cloudera Manager Agent incorrectly detected its own cgroup path

The Cloudera Manager Agent incorrectly detected its own cgroup path and created the YARN NodeManager's cgroup (for example, hadoop-yarn) under the Cloudera Manager Agent's system.slice hierarchy instead of the root-level cgroup path.

For example, the YARN cgroup appeared as /sys/fs/cgroup/cpu,cpuacct/system.slice/cloudera-scm-agent.service/hadoop-yarn instead of /sys/fs/cgroup/cpu,cpuacct/hadoop-yarn.

Because of this misplacement, when you restart the Cloudera Manager Agent process, systemd automatically destroys the nested cgroup (/system.slice/cloudera-scm-agent.service/hadoop-yarn). This immediately kills all running YARN containers and causes active jobs to fail.

To prevent YARN from inheriting the Cloudera Manager Agent's cgroup hierarchy, you can explicitly configure Cloudera Manager Agent to use the root cgroup path. Perform this configuration by uncommenting and setting the cgroups paths in the Cloudera Manager Agent configuration file: /etc/cloudera-scm-agent/config.ini. Perform the following steps:

1. Under the [cgroups] section, uncomment or add the following lines (adjusting for your cgroup version and controller types):

```
[cgroups]
mounts=cpu,cpuacct,cpuset,memory
cpu_cgroup_mount_point=/sys/fs/cgroup/cpu,cpuacct
memory_cgroup_mount_point=/sys/fs/cgroup/memory
```

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

```
sudo systemctl restart cloudera-scm-agent
```

This configuration ensures that the Cloudera Manager Agent and its managed roles (such as YARN NodeManager) always use root-level cgroup paths rather than inheriting them from system.slice, and prevents systemd from automatically cleaning up those cgroups when Cloudera Manager Agent restarts.

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

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.

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

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

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-73211: Cloudera Manager 7.11.3 does not clean up Python Path impacting Hue to start

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.11.3 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

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-72984: Alerts due to change in hostname fetching functionality in jdk 8 and jdk 11

Upgrading JAVA from JDK 8 to JDK 11 creates the following alert in CMS:

Bad : CMSERVER:pit666.slayer.mayank: Reaching Cloudera Manager Server failed

This happens due to a functionality change in JDK 11 on hostname fetching.

```
[root@pit666.slayer ~]# /us/lib/jvm/java-1.8.0/bin/java GetHostN
ame
Hostname: pit666.slayer.mayank

[root@pit666.slayer ~]# /usr/lib/jvm/java-11/bin/java GetHostName
Hostname: pit666.slayer
```

You can notice that the "hostname" is set to a short name instead of FQDN.

The current workaround is to set the hostname as FQDN.

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 7.11.3 or Cloudera Manager 7.11.3 CHF-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-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.

Manually add = sign to file_metadata_reload_properties configuration and modify the Gflags configuration file when the file_metadata_reload_properties configuration is passed as empty.

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.

To resolve this issue, disable the Enable Log Event Capture configuration on the Configuration page of a service.

OPSAPS-68845: Cloudera Manager Server fails to start after the Cloudera Manager upgrade

Starting from the Cloudera Manager 7.11.3 version up to the Cloudera Manager 7.11.3 CHF7 version, the Cloudera Manager Server fails to start after the Cloudera Manager upgrade due to Navigator user roles improperly handled in the upgrade in some scenarios.

None

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.

Upgrade to CDP 7.1.9, or downgrade to Cloudera Manager 7.7.1.

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

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

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

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

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

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

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

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

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

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

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

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

Azul Open JDK 8

RHEL or SLES

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

Ubuntu or Debian

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

Azul Open JDK 11

For DEBs only

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

OPSAPS-70207: Cloudera Manager Agents sending the Impala profile data with an incorrect header

Cloudera Manager agent might send incorrect HTTP header to Telemetry Publisher causing incorrect Content-Type error message resulting connection error. This issue causes missing Impala profile on Observatory.

Impala profile data is not available on Observatory.

Telemetry Publisher logs show:

DEBUG org.apache.cxf.jaxrs.utils.JAXRSUtils: No method match, method name : addProfileEvent, request path : /cluster/impala2, method @Path : /{clusterName}/{serviceName}, HTTP Method : POST, method HTTP Method : POST, ContentType : application/x-www-form-urlencoded, method @Consumes : application/json,, Accept : */*,, method @Produces : application/json,.

Cloudera Manager agent logs on Impalad hosts report:

Error occurred when sending entry to server: HTTP Error 415: Unsupported Media Type, url: http://<telemetry_publisher_host>:<port>

None

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

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

Upgrade the source Cloudera Manager to 7.11.3 or higher version.

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

The status of the deleted table displays as ACTIVE.

None

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

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

None

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Table 14: Cloudera Manager 7.11.3-CHF5

Repository Type	Repository Location

RHEL 9 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.7/redhat9/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.7/redhat9/yum/cloudera-manager.repo`

RHEL 8 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.7/redhat8/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.7/redhat8/yum/cloudera-manager.repo`

RHEL 7 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.7/redhat7/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.7/redhat7/yum/cloudera-manager.repo`

SLES 15	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.7/sles15/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.7/sles15/yum/cloudera-manager.repo`

SLES 12	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.7/sles12/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.7/sles12/yum/cloudera-manager.repo`

Repository Type	Repository Location

Ubuntu 20	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.7/ubuntu2004/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.7/ubuntu2004/apt/cloudera-manager.list`

IBM PowerPC RHEL 7	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.7/redhat7-ppc/yum`
IBM PowerPC RHEL 8	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.7/redhat8-ppc/yum`

Cloudera Manager 7.11.3 Cumulative hotfix 4

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 4.

This cumulative hotfix was released on March 8, 2024.

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

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

Added FIPS support for JDK11 in Zeppelin.

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF4 (version: 7.11.3.6-50817646):
OPSAPS-75899: HDFS directory creation fails on JDK 11 or higher when LDAP or Active Directory integrated clusters using the hadoop.security.group.mapping property.

Due to this issue, many critical Cloudera Manager operations fail to complete, such as:

- Install Oozie ShareLib (Oozie Actions Install Oozie ShareLib)
- Install YARN MapReduce Framework JARs (YARN Actions Install YARN MapReduce Framework JARs)

You must perform the following workaround steps to manually add specific Java modules to the HDFS service script on all nodes in the cluster:

1. Navigate to the directory /opt/cloudera/cm-agent/service/hdfs/.
2. Open the hdfs.sh file for editing.
3. Locate the line starting with MKDIR_JAVA_OPTS.
4. Update it to include the following flags:

```
Change this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=f
ile://${MKDIR_LOG4J_FILE}"
To this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=fi
le://${MKDIR_LOG4J_FILE} --add-exports=java.naming/com.sun.j
```

```
ndi.ldap=ALL-UNNAMED --add-opens=java.naming/com.sun.jndi.ld
ap=ALL-UNNAMED"
```

OPSAPS-70915: Cloudera Manager Agent incorrectly detected its own cgroup path

The Cloudera Manager Agent incorrectly detected its own cgroup path and created the YARN NodeManager's cgroup (for example, hadoop-yarn) under the Cloudera Manager Agent's system.slice hierarchy instead of the root-level cgroup path.

For example, the YARN cgroup appeared as /sys/fs/cgroup/cpu,cpuacct/system.slice/cloudera-scm-agent.service/hadoop-yarn instead of /sys/fs/cgroup/cpu,cpuacct/hadoop-yarn.

Because of this misplacement, when you restart the Cloudera Manager Agent process, systemd automatically destroys the nested cgroup (/system.slice/cloudera-scm-agent.service/hadoop-yarn). This immediately kills all running YARN containers and causes active jobs to fail.

To prevent YARN from inheriting the Cloudera Manager Agent's cgroup hierarchy, you can explicitly configure Cloudera Manager Agent to use the root cgroup path. Perform this configuration by uncommenting and setting the cgroups paths in the Cloudera Manager Agent configuration file: /etc/cloudera-scm-agent/config.ini. Perform the following steps:

1. Under the [cgroups] section, uncomment or add the following lines (adjusting for your cgroup version and controller types):

   ```
   [cgroups]
   mounts=cpu,cpuacct,cpuset,memory
   cpu_cgroup_mount_point=/sys/fs/cgroup/cpu,cpuacct
   memory_cgroup_mount_point=/sys/fs/cgroup/memory
   ```

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

   ```
   sudo systemctl restart cloudera-scm-agent
   ```

   This configuration ensures that the Cloudera Manager Agent and its managed roles (such as YARN NodeManager) always use root-level cgroup paths rather than inheriting them from syst em.slice, and prevents systemd from automatically cleaning up those cgroups when Cloudera Manager Agent restarts.

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

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.

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

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

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-73211: Cloudera Manager 7.11.3 does not clean up Python Path impacting Hue to start

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.11.3 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

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-72984: Alerts due to change in hostname fetching functionality in jdk 8 and jdk 11

Upgrading JAVA from JDK 8 to JDK 11 creates the following alert in CMS:

Bad : CMSERVER:pit666.slayer.mayank: Reaching Cloudera Manager Server failed

This happens due to a functionality change in JDK 11 on hostname fetching.

```
[root@pit666.slayer ~]# /us/lib/jvm/java-1.8.0/bin/java GetHostN
ame
Hostname: pit666.slayer.mayank

[root@pit666.slayer ~]# /usr/lib/jvm/java-11/bin/java GetHostName
Hostname: pit666.slayer
```

You can notice that the "hostname" is set to a short name instead of FQDN.

The current workaround is to set the hostname as FQDN.

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 7.11.3 or Cloudera Manager 7.11.3 CHF-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-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.

Manually add = sign to file_metadata_reload_properties configuration and modify the Gflags configuration file when the file_metadata_reload_properties configuration is passed as empty.

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.

To resolve this issue, disable the Enable Log Event Capture configuration on the Configuration page of a service.

OPSAPS-68845: Cloudera Manager Server fails to start after the Cloudera Manager upgrade

Starting from the Cloudera Manager 7.11.3 version up to the Cloudera Manager 7.11.3 CHF7 version, the Cloudera Manager Server fails to start after the Cloudera Manager upgrade due to Navigator user roles improperly handled in the upgrade in some scenarios.

None

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.

Upgrade to CDP 7.1.9, or downgrade to Cloudera Manager 7.7.1.

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

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

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

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

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

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

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

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

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

Azul Open JDK 8

RHEL or SLES

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

Ubuntu or Debian

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

Azul Open JDK 11

For DEBs only

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

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

The status of the deleted table displays as ACTIVE.

None

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

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

None

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

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

Upgrade the source Cloudera Manager to 7.11.3 or higher version.

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

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

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

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

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

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

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

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

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

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

Cloudera Manager now registers kafka_connect_connector_task_metrics_batch_size_avg and kafka_connect_connector_task_metrics_batch_size_max metrics correctly.

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

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

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

The decommission now works as expected.

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

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

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

Table 15: Cloudera Manager 7.11.3-CHF4

Repository Type	Repository Location

RHEL 9 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.6/redhat9/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.6/redhat9/yum/cloudera-manager.repo`

RHEL 8 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.6/redhat8/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.6/redhat8/yum/cloudera-manager.repo`

RHEL 7 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.6/redhat7/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.6/redhat7/yum/cloudera-manager.repo`

SLES 15	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.6/sles15/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.6/sles15/yum/cloudera-manager.repo`

SLES 12	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.6/sles12/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.6/sles12/yum/cloudera-manager.repo`

Repository Type	Repository Location

Ubuntu 20	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.6/ubuntu2004/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/7.11.3.6/ubuntu2004/apt/cloudera-manager.list`

Cloudera Manager 7.11.3 Cumulative hotfix 3

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 3.

This cumulative hotfix was released on February 23, 2024.

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

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF3 (version: 7.11.3.4-50275000):
OPSAPS-75899: HDFS directory creation fails on JDK 11 or higher when LDAP or Active Directory integrated clusters using the hadoop.security.group.mapping property.

Due to this issue, many critical Cloudera Manager operations fail to complete, such as:

- Install Oozie ShareLib (Oozie Actions Install Oozie ShareLib)
- Install YARN MapReduce Framework JARs (YARN Actions Install YARN MapReduce Framework JARs)

You must perform the following workaround steps to manually add specific Java modules to the HDFS service script on all nodes in the cluster:

1. Navigate to the directory /opt/cloudera/cm-agent/service/hdfs/.
2. Open the hdfs.sh file for editing.
3. Locate the line starting with MKDIR_JAVA_OPTS.
4. Update it to include the following flags:

```
Change this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=f
ile://${MKDIR_LOG4J_FILE}"
To this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=fi
le://${MKDIR_LOG4J_FILE} --add-exports=java.naming/com.sun.j
ndi.ldap=ALL-UNNAMED --add-opens=java.naming/com.sun.jndi.ld
ap=ALL-UNNAMED"
```

OPSAPS-70915: Cloudera Manager Agent incorrectly detected its own cgroup path

The Cloudera Manager Agent incorrectly detected its own cgroup path and created the YARN NodeManager's cgroup (for example, hadoop-yarn) under the Cloudera Manager Agent's system.slice hierarchy instead of the root-level cgroup path.

For example, the YARN cgroup appeared as /sys/fs/cgroup/cpu,cpuacct/system.slice/cloudera-scm-agent.service/hadoop-yarn instead of /sys/fs/cgroup/cpu,cpuacct/hadoop-yarn.

Because of this misplacement, when you restart the Cloudera Manager Agent process, systemd automatically destroys the nested cgroup (/system.slice/cloudera-scm-agent.service/hadoop-yarn). This immediately kills all running YARN containers and causes active jobs to fail.

To prevent YARN from inheriting the Cloudera Manager Agent's cgroup hierarchy, you can explicitly configure Cloudera Manager Agent to use the root cgroup path. Perform this configuration by uncommenting and setting the cgroups paths in the Cloudera Manager Agent configuration file: /etc/cloudera-scm-agent/config.ini. Perform the following steps:

1. Under the [cgroups] section, uncomment or add the following lines (adjusting for your cgroup version and controller types):

```
[cgroups]
mounts=cpu,cpuacct,cpuset,memory
cpu_cgroup_mount_point=/sys/fs/cgroup/cpu,cpuacct
memory_cgroup_mount_point=/sys/fs/cgroup/memory
```

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

```
sudo systemctl restart cloudera-scm-agent
```

This configuration ensures that the Cloudera Manager Agent and its managed roles (such as YARN NodeManager) always use root-level cgroup paths rather than inheriting them from syst em.slice, and prevents systemd from automatically cleaning up those cgroups when Cloudera Manager Agent restarts.

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

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.

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

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

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-73211: Cloudera Manager 7.11.3 does not clean up Python Path impacting Hue to start

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.11.3 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

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-72984: Alerts due to change in hostname fetching functionality in jdk 8 and jdk 11

Upgrading JAVA from JDK 8 to JDK 11 creates the following alert in CMS:

Bad : CMSERVER:pit666.slayer.mayank: Reaching Cloudera Manager Server failed

This happens due to a functionality change in JDK 11 on hostname fetching.

```
[root@pit666.slayer ~]# /us/lib/jvm/java-1.8.0/bin/java GetHostN
ame
Hostname: pit666.slayer.mayank

[root@pit666.slayer ~]# /usr/lib/jvm/java-11/bin/java GetHostName
Hostname: pit666.slayer
```

You can notice that the "hostname" is set to a short name instead of FQDN.

The current workaround is to set the hostname as FQDN.

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 7.11.3 or Cloudera Manager 7.11.3 CHF-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-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.

Manually add = sign to file_metadata_reload_properties configuration and modify the Gflags configuration file when the file_metadata_reload_properties configuration is passed as empty.

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.

To resolve this issue, disable the Enable Log Event Capture configuration on the Configuration page of a service.

OPSAPS-68845: Cloudera Manager Server fails to start after the Cloudera Manager upgrade

Starting from the Cloudera Manager 7.11.3 version up to the Cloudera Manager 7.11.3 CHF7 version, the Cloudera Manager Server fails to start after the Cloudera Manager upgrade due to Navigator user roles improperly handled in the upgrade in some scenarios.

None

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.

Upgrade to CDP 7.1.9, or downgrade to Cloudera Manager 7.7.1.

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

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

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

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

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

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

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

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

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

Azul Open JDK 8

RHEL or SLES

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

Ubuntu or Debian

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

Azul Open JDK 11

For DEBs only

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

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

The status of the deleted table displays as ACTIVE.

None

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

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

None

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

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

Upgrade the source Cloudera Manager to 7.11.3 or higher version.

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

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

Contact Cloudera support for a workaround.

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

When Cloudera Manager is upgraded to 7.11.3, if the CDP cluster is not 7.1.9, then the YARN Container Usage Aggregation job fails.

Add the following property in the MapReduce Client Advanced Configuration Snippet (Safety Valve) for mapred-site.xml file.

```
NAME: mapreduce.jvm.add-opens-as-default
VALUE: false
```

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

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

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

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

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

Cloudera Manager replaces the '--' characters in XML configuration file comments with &#8212 which is the Unicode character of '--'.

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

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

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

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

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

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

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

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

OPSAPS-69063: Concurrent policy creation to multiple targets

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

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

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

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

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

OPSAPS-69194: JDK 17 support for HBase Indexer

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

OPSAPS-69378: Increase default certificate lifetime

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

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

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

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

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

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

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

OPSAPS-69245: Oozie issue in FIPS environments

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

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

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

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

Table 16: Cloudera Manager 7.11.3-CHF3

Repository Type	Repository Location

RHEL 9 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.4-50275000/redhat9/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.4-50275000/redhat9/yum/cloudera-manager.repo`

Repository Type	Repository Location

RHEL 8 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.4-50275000/redhat8/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.4-50275000/redhat8/yum/cloudera-manager.repo`

RHEL 7 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.4-50275000/redhat7/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.4-50275000/redhat7/yum/cloudera-manager.repo`

SLES 15	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.4-50275000/sles15/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.4-50275000/sles15/yum/cloudera-manager.repo`

SLES 12	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.4-50275000/sles12/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.4-50275000/sles12/yum/cloudera-manager.repo`

Ubuntu 20	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.4-50275000/ubuntu2004/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.4-50275000/ubuntu2004/apt/cloudera-manager.list`

Cloudera Manager 7.11.3 Cumulative hotfix 2

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 2.

This cumulative hotfix was released on December 21, 2023.

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

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

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

Wait timeout for regenerating credentials in Active Directory (AD)

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

Important:

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

You might set this parameter before running the following commands:

- Generate Missing Credentials
- Regenerate Selected

FIPS support for JDK11 in Kudu

Added FIPS support for JDK11 in Kudu.

FIPS support for JDK11 in Hive

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

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF2 (version: 7.11.3.3-47960007):
OPSAPS-75899: HDFS directory creation fails on JDK 11 or higher when LDAP or Active Directory integrated clusters using the hadoop.security.group.mapping property.

Due to this issue, many critical Cloudera Manager operations fail to complete, such as:

- Install Oozie ShareLib (Oozie Actions Install Oozie ShareLib)
- Install YARN MapReduce Framework JARs (YARN Actions Install YARN MapReduce Framework JARs)

You must perform the following workaround steps to manually add specific Java modules to the HDFS service script on all nodes in the cluster:

1. Navigate to the directory /opt/cloudera/cm-agent/service/hdfs/.
2. Open the hdfs.sh file for editing.
3. Locate the line starting with MKDIR_JAVA_OPTS.
4. Update it to include the following flags:

```
Change this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=f
ile://${MKDIR_LOG4J_FILE}"
To this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=fi
le://${MKDIR_LOG4J_FILE} --add-exports=java.naming/com.sun.j
```

```
ndi.ldap=ALL-UNNAMED --add-opens=java.naming/com.sun.jndi.ld
ap=ALL-UNNAMED"
```

OPSAPS-70915: Cloudera Manager Agent incorrectly detected its own cgroup path

The Cloudera Manager Agent incorrectly detected its own cgroup path and created the YARN NodeManager's cgroup (for example, hadoop-yarn) under the Cloudera Manager Agent's system.slice hierarchy instead of the root-level cgroup path.

For example, the YARN cgroup appeared as /sys/fs/cgroup/cpu,cpuacct/system.slice/cloudera-scm-agent.service/hadoop-yarn instead of /sys/fs/cgroup/cpu,cpuacct/hadoop-yarn.

Because of this misplacement, when you restart the Cloudera Manager Agent process, systemd automatically destroys the nested cgroup (/system.slice/cloudera-scm-agent.service/hadoop-yarn). This immediately kills all running YARN containers and causes active jobs to fail.

To prevent YARN from inheriting the Cloudera Manager Agent's cgroup hierarchy, you can explicitly configure Cloudera Manager Agent to use the root cgroup path. Perform this configuration by uncommenting and setting the cgroups paths in the Cloudera Manager Agent configuration file: /etc/cloudera-scm-agent/config.ini. Perform the following steps:

1. Under the [cgroups] section, uncomment or add the following lines (adjusting for your cgroup version and controller types):

   ```
   [cgroups]
   mounts=cpu,cpuacct,cpuset,memory
   cpu_cgroup_mount_point=/sys/fs/cgroup/cpu,cpuacct
   memory_cgroup_mount_point=/sys/fs/cgroup/memory
   ```

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

   ```
   sudo systemctl restart cloudera-scm-agent
   ```

   This configuration ensures that the Cloudera Manager Agent and its managed roles (such as YARN NodeManager) always use root-level cgroup paths rather than inheriting them from syst em.slice, and prevents systemd from automatically cleaning up those cgroups when Cloudera Manager Agent restarts.

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

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.

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

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

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-73211: Cloudera Manager 7.11.3 does not clean up Python Path impacting Hue to start

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.11.3 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

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-72984: Alerts due to change in hostname fetching functionality in jdk 8 and jdk 11

Upgrading JAVA from JDK 8 to JDK 11 creates the following alert in CMS:

Bad : CMSERVER:pit666.slayer.mayank: Reaching Cloudera Manager Server failed

This happens due to a functionality change in JDK 11 on hostname fetching.

```
[root@pit666.slayer ~]# /us/lib/jvm/java-1.8.0/bin/java GetHostN
ame
Hostname: pit666.slayer.mayank

[root@pit666.slayer ~]# /usr/lib/jvm/java-11/bin/java GetHostName
Hostname: pit666.slayer
```

You can notice that the "hostname" is set to a short name instead of FQDN.

The current workaround is to set the hostname as FQDN.

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 7.11.3 or Cloudera Manager 7.11.3 CHF-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-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.

Manually add = sign to file_metadata_reload_properties configuration and modify the Gflags configuration file when the file_metadata_reload_properties configuration is passed as empty.

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.

Upgrade to CDP 7.1.9, or downgrade to Cloudera Manager 7.7.1.

OPSAPS-68845: Cloudera Manager Server fails to start after the Cloudera Manager upgrade

Starting from the Cloudera Manager 7.11.3 version up to the Cloudera Manager 7.11.3 CHF7 version, the Cloudera Manager Server fails to start after the Cloudera Manager upgrade due to Navigator user roles improperly handled in the upgrade in some scenarios.

None

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

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

None.

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

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

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

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

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

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

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

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

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

Azul Open JDK 8

RHEL or SLES

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

Ubuntu or Debian

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

Azul Open JDK 11

For DEBs only

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

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

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

Perform the following steps:

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

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

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

```
sudo systemctl restart cloudera-scm-server
```

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

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

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

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

The status of the deleted table displays as ACTIVE.

None

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

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

None

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

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

Upgrade the source Cloudera Manager to 7.11.3 or higher version.

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

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

Contact Cloudera support for a workaround.

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

When Cloudera Manager is upgraded to 7.11.3, if the CDP cluster is not 7.1.9, then the YARN Container Usage Aggregation job fails.

Add the following property in the MapReduce Client Advanced Configuration Snippet (Safety Valve) for mapred-site.xml file.

```
NAME: mapreduce.jvm.add-opens-as-default
VALUE: false
```

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

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

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

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

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

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

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

OPSAPS-68722: Java heap size can now be configured

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

OPSAPS-68217: Add post replication diff to compare files

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

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

      SCHEDULES_WITH_ADDITIONAL_DEBUG_STEPS = [***comma-separated list of numerical IDs of all the applicable replication policies***]
   c. Search for the HDFS Replication Advanced Configuration Snippet (Safety Valve) for hdfs-site.xml property.
   d. Add the following key-value pair, and Save the changes:

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

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

      

      Note: The key-value pair fails the replication job if there are any missing files on the target. This failure does not invalidate the snapshot, which means that if a file is missing, it remains missing until you manually force-run a bootstrap replication which is an optional action. If you do not set the key-value pair, the replication job continues and generates the debug info.
3. To enable extra logging, complete the following steps:
   a. Search for the HDFS Replication Environment Advanced Configuration Snippet (Safety Valve) property.
   b. Add the following key-value pair, and Save the changes:

      EXTRA_LOG_CONFIGS_$SCHEDULE_ID =

```
log4j.rootLogger=INFO,console;
hadoop.root.logger=INFO,console;log4j.appender.console=or
g.apache.log4j.ConsoleAppender;
log4j.appender.console.target=System.err;log4j.appender.cons
ole.layout=org.apache.log4j.PatternLayout;
```

```
log4j.appender.console.layout.ConversionPattern=%d{yy/MM/dd
 HH:mm:ss} %p %c{2}: %m%n;
log4j.logger.org.apache.hadoop.fs.azurebfs.services.AbfsIoU
tils=DEBUG,console;
log4j.logger.org.apache.hadoop.fs.azurebfs.services.AbfsClie
nt=DEBUG,console;
log4j.logger.distcp.SimpleCopyListing=DEBUG,console;log4j.
logger.distcp.SnapshotDiffGenerator=DEBUG,console
```

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

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

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

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

OPSAPS-68516: Ozone replication diagnostic bundle collection

Replication Manager generates diagnostic information bundle for Ozone replication policies.

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

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

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

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

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

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

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

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

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

OPSAPS-68526 - Iceberg Replication support for Dell Powerscale

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

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

Table 17: Cloudera Manager 7.11.3-CHF2

Repository Type	Repository Location

RHEL 9 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.3-47960007/redhat9/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.3-47960007/redhat9/yum/cloudera-manager.repo`

RHEL 8 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.3-47960007/redhat8/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.3-47960007/redhat8/yum/cloudera-manager.repo`

RHEL 7 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.3-47960007/redhat7/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.3-47960007/redhat7/yum/cloudera-manager.repo`

SLES 15	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.3-47960007/sles15/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.3-47960007/sles15/yum/cloudera-manager.repo`

SLES 12	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.3-47960007/sles12/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.3-47960007/sles12/yum/cloudera-manager.repo`

Repository Type	Repository Location

Ubuntu 20	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.3-47960007/ubuntu2004/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.3-47960007/ubuntu2004/apt/cloudera-manager.list`

IBM PowerPC RHEL 9	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.3-47960007/redhat9-ppc/yum`
IBM PowerPC RHEL 8	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.3-47960007/redhat8-ppc/yum`

Cloudera Manager 7.11.3 Cumulative hotfix 1

Know more about the Cloudera Manager 7.11.3 cumulative hotfixes 1.

This cumulative hotfix was released on November 2, 2023.

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

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

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

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

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

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

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

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

Following are the list of known issues and their corresponding workarounds that are shipped for Cloudera Manager 7.11.3 CHF1 (version: 7.11.3.2-46642574):
OPSAPS-75899: HDFS directory creation fails on JDK 11 or higher when LDAP or Active Directory integrated clusters using the hadoop.security.group.mapping property.

Due to this issue, many critical Cloudera Manager operations fail to complete, such as:

- Install Oozie ShareLib (Oozie Actions Install Oozie ShareLib)

- Install YARN MapReduce Framework JARs (YARN Actions Install YARN MapReduce Framework JARs)

You must perform the following workaround steps to manually add specific Java modules to the HDFS service script on all nodes in the cluster:

1. Navigate to the directory /opt/cloudera/cm-agent/service/hdfs/.
2. Open the hdfs.sh file for editing.
3. Locate the line starting with MKDIR_JAVA_OPTS.
4. Update it to include the following flags:

```
Change this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=f
ile://${MKDIR_LOG4J_FILE}"
To this:
MKDIR_JAVA_OPTS="${MKDIR_JAVA_OPTS} -Dlog4j.configuration=fi
le://${MKDIR_LOG4J_FILE} --add-exports=java.naming/com.sun.j
ndi.ldap=ALL-UNNAMED --add-opens=java.naming/com.sun.jndi.ld
ap=ALL-UNNAMED"
```

OPSAPS-70915: Cloudera Manager Agent incorrectly detected its own cgroup path

The Cloudera Manager Agent incorrectly detected its own cgroup path and created the YARN NodeManager's cgroup (for example, hadoop-yarn) under the Cloudera Manager Agent's system.slice hierarchy instead of the root-level cgroup path.

For example, the YARN cgroup appeared as /sys/fs/cgroup/cpu,cpuacct/system.slice/cloudera-scm-agent.service/hadoop-yarn instead of /sys/fs/cgroup/cpu,cpuacct/hadoop-yarn.

Because of this misplacement, when you restart the Cloudera Manager Agent process, systemd automatically destroys the nested cgroup (/system.slice/cloudera-scm-agent.service/hadoop-yarn). This immediately kills all running YARN containers and causes active jobs to fail.

To prevent YARN from inheriting the Cloudera Manager Agent's cgroup hierarchy, you can explicitly configure Cloudera Manager Agent to use the root cgroup path. Perform this configuration by uncommenting and setting the cgroups paths in the Cloudera Manager Agent configuration file: /etc/cloudera-scm-agent/config.ini. Perform the following steps:

1. Under the [cgroups] section, uncomment or add the following lines (adjusting for your cgroup version and controller types):

```
[cgroups]
mounts=cpu,cpuacct,cpuset,memory
cpu_cgroup_mount_point=/sys/fs/cgroup/cpu,cpuacct
memory_cgroup_mount_point=/sys/fs/cgroup/memory
```

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

```
sudo systemctl restart cloudera-scm-agent
```

This configuration ensures that the Cloudera Manager Agent and its managed roles (such as YARN NodeManager) always use root-level cgroup paths rather than inheriting them from system.slice, and prevents systemd from automatically cleaning up those cgroups when Cloudera Manager Agent restarts.

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

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.

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

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

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-73211: Cloudera Manager 7.11.3 does not clean up Python Path impacting Hue to start

When you upgrade from Cloudera Manager 7.7.1 or lower versions to Cloudera Manager 7.11.3 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

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-72984: Alerts due to change in hostname fetching functionality in jdk 8 and jdk 11

Upgrading JAVA from JDK 8 to JDK 11 creates the following alert in CMS:

Bad : CMSERVER:pit666.slayer.mayank: Reaching Cloudera Manager Server failed

This happens due to a functionality change in JDK 11 on hostname fetching.

```
[root@pit666.slayer ~]# /us/lib/jvm/java-1.8.0/bin/java GetHostN
ame
Hostname: pit666.slayer.mayank

[root@pit666.slayer ~]# /usr/lib/jvm/java-11/bin/java GetHostName
Hostname: pit666.slayer
```

You can notice that the "hostname" is set to a short name instead of FQDN.

The current workaround is to set the hostname as FQDN.

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 7.11.3 or Cloudera Manager 7.11.3 CHF-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-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.

Manually add = sign to file_metadata_reload_properties configuration and modify the Gflags configuration file when the file_metadata_reload_properties configuration is passed as empty.

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.

Upgrade to CDP 7.1.9, or downgrade to Cloudera Manager 7.7.1.

OPSAPS-68845: Cloudera Manager Server fails to start after the Cloudera Manager upgrade

Starting from the Cloudera Manager 7.11.3 version up to the Cloudera Manager 7.11.3 CHF7 version, the Cloudera Manager Server fails to start after the Cloudera Manager upgrade due to Navigator user roles improperly handled in the upgrade in some scenarios.

None

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

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

None.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Azul Open JDK 8

RHEL or SLES

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

Ubuntu or Debian

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

Azul Open JDK 11

For DEBs only

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

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

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

Contact Cloudera support for a workaround.

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

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

None

OPSAPS-68658: Source ozone service id used as target

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

None

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

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

None

OPSAPS-42908: "User:hdfs not allowed to do DECRYPT_EEK" error appears for Hive external table replication policies

When you run Hive external table replication policies on clusters using Ranger KMS, the “User:hdfs not allowed to do 'DECRYPT_EEK'” error appears when you do not use the hdfs username.

Edit the Hive external table replication policy, and configure the Advanced Directory for metadata file field to a new directory that is not encrypted. The replication policy uses this directory to store the transient data during Hive replication.

OPSAPS-69897: NPE in Ozone replication from CM 7.7.1 to CM 7.11.3

When you use source Cloudera Manager 7.7.1 and target Cloudera Manager 7.11.3 for Ozone replication policies, the policies fail with Failure during PreOzoneCopyListingCheck execution: null error. This is because the target Cloudera Manager 7.11.3 does not retrieve the required source bucket information for validation from the source Cloudera Manager 7.7.1 during the PreCopyListingCheck command phase. You come across this error when you use source Cloudera Manager versions lower than 7.10.1 and target Cloudera Manager versions higher than or equal to 7.10.1 in an Ozone replication policy.

Upgrade the source Cloudera Manager to 7.11.3 or higher version.

CDPD-62464: Java process called by navatlas.sh tool fails on JDK-8 version

While running nav2atlas.sh script on OracleJDK 8 an error message is thrown and returns code 0 on an unsuccessful run.

You must install JDK-11 version on the host. Make sure not to put into the default path and JAVA _HOME. In a shell, set the JAVA_HOME to this location and run the nav2atlas.sh script.

CDPD-62834: Status of the deleted table is seen as ACTIVE in Atlas after the completion of navigator2atlas migration process

The status of the deleted table displays as ACTIVE.

None

CDPD-62837: During the navigator2atlas process, the hive_storagedesc is incomplete in Atlas

For the hive_storagedesc entity, some of the attributes are not getting populated.

None

Following are the list of fixed issues that were shipped for Cloudera Manager 7.11.3 CHF1 (version: 7.11.3.2-46642574):

OPSAPS-68664: Added the support for JDK 17 in HDFS.

This issue is resolved.

OPSAPS-68550: Ozone Canary failing with unknown option --skipTrash.

This issue is resolved.

OPSAPS-66023: Error message about an unsupported ciphersuite while upgrading or installing cluster with the latest FIPS compliance

When upgrading or installing a FIPS enabled cluster, Cloudera Manager is unable to download the new CDP parcel from the Cloudera parcel archive.

Cloudera Manager displays the following error message:

HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA

This issue is fixed now by correcting the incorrect ciphersuite selection.

OPSAPS-65504: Upgraded Apache Ivy version

The Apache Ivy version is upgraded from 2.x.x to 2.5.1 version to fix CVE issues.

OPSAPS-68500: The cloudera-manager-installer.bin fails to reach Ubuntu 20 repository on the Archive URL due to redirections

Agent Installation with Cloudera Manager on Ubuntu20 platform does not function when the self-installer method (using the installer.bin file) is employed to install Cloudera Manager. The failure mode is that Cloudera Manager Agent installation step will fail with an error message saying "The

repository 'https://archive.cloudera.com/p/cm7/7.11.3/ubuntu2004/apt focal-cm7 InRelease' is not signed."

This issue is fixed now.

OPSAPS-68422: Incorrect HBase shutdown command can lead to inconsistencies

Cloudera Manager uses an incomplete stop command when you stop the HBase service or the corresponding roles on a 7.1.8 or higher private cloud cluster. Due to this, the Cloudera Manager cannot gracefully stop the processes and kill them after a set timeout. This could lead to metadata corruption.

This issue is fixed now.

OPSAPS-68506: Knox CSD changes for readiness check

A readiness endpoint was added to determine whether Knox is ready to receive traffic. Clouder Manager checks the state of Knox after startup to reduce downtime during rolling restarts.

OPSAPS-68424 Impala: CM Agent unable to extract logs to TP export directory

Impala queries were not displaying in the Cloudera Observability and Workload XM web User Interfaces. This was due to an internal error that was stopping Cloudera Manager from pushing the Impala profile to the Telemetry Publisher logs directory.

This Issue is now fixed and the Telemetry Publisher's log extraction has been re-enable.

Note: If you are using a Cloudera Manager version between 7.11.2.0 and 7.11.3, Cloudera recommends upgrading to Cloudera Manager 7.11.3.CHF1.

OPSAPS-68697 Error while generating email template resulting in an inability to trigger mail notification

Cloudera Observability and Workload XM were unable to trigger an email notification when an Impala query matched the Auto Action's alert threshold value.

This problem occurred when both the following conditions were met:

- The Auto Action is triggered for an Impala Scope.
- The length of the Impala query on which the action event is triggered is less than 36 characters.

This issue is now fixed.

OPSAPS-69480: Hardcode MR add-opens-as-default config

When Cloudera Manager is upgraded to 7.11.3, if the CDP cluster is not 7.1.9, then the YARN Container Usage Aggregation job fails.

Add the following property in the MapReduce Client Advanced Configuration Snippet (Safety Valve) for mapred-site.xml file.

```
NAME: mapreduce.jvm.add-opens-as-default
VALUE: false
```

OPSAPS-68798 Auto Actions not using proxy while connecting to DBUS

The Cloudera Observability and Workload XM Auto Actions feature was not recognizing the proxy server credentials, even when they were correct and the proxy server was enabled in Telemetry Publisher.

This issue is now fixed.

Note: Users must ensure they have the correct proxy server setup and that the proxy server is enabled in Telemetry Publisher.

Known issue:

OPSAPS-68629: HDFS HTTPFS GateWay is not able to start with custom krb5.conf location set in Cloudera Manager.

On a cluster with a custom krb5.conf file location configured in Cloudera Manager, HDFS HTTPFS role is not able to start because it does not have the custom Kerberos configuration file setting properly propagated to the service, and therefore it fails with a Kerberos related exception: in threa d "main" java.io.IOException: Unable to initialize WebAppContext at org.apache.hadoop.http .HttpServer2.start(HttpServer2.java:1240) at org.apache.hadoop.fs.http.server.HttpFSServerWebS erver.start(HttpFSServerWebServer.java:131) at org.apache.hadoop.fs.http.server.HttpFSServer WebServer.main(HttpFSServerWebServer.java:162) Caused by: java.lang.IllegalArgumentExc eption: Can't get Kerberos realm at org.apache.hadoop.security.HadoopKerberosName.setConfi guration(HadoopKerberosName.java:71) at org.apache.hadoop.security.UserGroupInformation. initialize(UserGroupInformation.java:329) at org.apache.hadoop.security.UserGroupInforma tion.setConfiguration(UserGroupInformation.java:380) at org.apache.hadoop.lib.service.hadoop .FileSystemAccessService.init(FileSystemAccessService.java:166) at org.apache.hadoop.lib.s erver.BaseService.init(BaseService.java:71) at org.apache.hadoop.lib.server.Server.initService s(Server.java:581) at org.apache.hadoop.lib.server.Server.init(Server.java:377) at org.apache.ha doop.fs.http.server.HttpFSServerWebApp.init(HttpFSServerWebApp.java:100) at org.apache.h adoop.lib.servlet.ServerWebApp.contextInitialized(ServerWebApp.java:158) at org.eclipse.jett y.server.handler.ContextHandler.callContextInitialized(ContextHandler.java:1073) at org.ecli pse.jetty.servlet.ServletContextHandler.callContextInitialized(ServletContextHandler.java:572) at org.eclipse.jetty.server.handler.ContextHandler.contextInitialized(ContextHandler.java:1002) at org.eclipse.jetty.servlet.ServletHandler.initialize(ServletHandler.java:765) at org.eclipse.jetty. servlet.ServletContextHandler.startContext(ServletContextHandler.java:379) at org.eclipse.jett y.webapp.WebAppContext.startWebapp(WebAppContext.java:1449) at org.eclipse.jetty.webap p.WebAppContext.startContext(WebAppContext.java:1414) at org.eclipse.jetty.server.handler. ContextHandler.doStart(ContextHandler.java:916) at org.eclipse.jetty.servlet.ServletContextHan dler.doStart(ServletContextHandler.java:288) at org.eclipse.jetty.webapp.WebAppContext.doStart (WebAppContext.java:524) at org.eclipse.jetty.util.component.AbstractLifeCycle.start(AbstractL ifeCycle.java:73) at org.eclipse.jetty.util.component.ContainerLifeCycle.start(ContainerLifeCycle. java:169) at org.eclipse.jetty.util.component.ContainerLifeCycle.doStart(ContainerLifeCycle. java:117) at org.eclipse.jetty.server.handler.AbstractHandler.doStart(AbstractHandler.java:97) a t org.eclipse.jetty.util.component.AbstractLifeCycle.start(AbstractLifeCycle.java:73) at org.ecl ipse.jetty.util.component.ContainerLifeCycle.start(ContainerLifeCycle.java:169) at org.eclipse.jet ty.server.Server.start(Server.java:423) at org.eclipse.jetty.util.component.ContainerLifeCycle.doS tart(ContainerLifeCycle.java:110) at org.eclipse.jetty.server.handler.AbstractHandler.doStart(Abst ractHandler.java:97) at org.eclipse.jetty.server.Server.doStart(Server.java:387) at org.eclipse.jett y.util.component.AbstractLifeCycle.start(AbstractLifeCycle.java:73) at org.apache.hadoop.http. HttpServer2.start(HttpServer2.java:1218) ... 2 more Caused by: java.lang.IllegalArgumentExcept ion: KrbException: Cannot locate default realm at java.security.jgss/javax.security.auth.kerberos. KerberosPrincipal.<init>(KerberosPrincipal.java:174) at org.apache.hadoop.security.authentic ation.util.KerberosUtil.getDefaultRealm(KerberosUtil.java:108) at org.apache.hadoop.security .HadoopKerberosName.setConfiguration(HadoopKerberosName.java:69) ...

1. Log in to Cloudera Manager.
2. Select the HDFS service.
3. Select Configurations tab.
4. Search for HttpFS Environment Advanced Configuration Snippet (Safety Valve)
5. Add to or extend the HADOOP_OPTS environment variable with the following value: -Djava.security.krb5.conf=<the custom krb5.conf location>
6. Click Save Changes.

The repositories for Cloudera Manager 7.11.3-CHF1 are listed in the following table:

Table 18: Cloudera Manager 7.11.3-CHF1

Repository Type	Repository Location

RHEL 9 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.2-46642574/redhat9/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.2-46642574/redhat9/yum/cloudera-manager.repo`

RHEL 8 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.2-46642574/redhat8/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.2-46642574/redhat8/yum/cloudera-manager.repo`

RHEL 7 Compatible	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.2-46642574/redhat7/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.2-46642574/redhat7/yum/cloudera-manager.repo`

SLES 15	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.2-46642574/sles15/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.2-46642574/sles15/yum/cloudera-manager.repo`

SLES 12	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.2-46642574/sles12/yum`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.2-46642574/sles12/yum/cloudera-manager.repo`

Repository Type	Repository Location

Ubuntu 20	Repository:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.2-46642574/ubuntu2004/apt`
	Repository File:
	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.2-46642574/ubuntu2004/apt/cloudera-manager.repo`

IBM PowerPC RHEL 8	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.2-46642574/redhat8-ppc/yum`
IBM PowerPC RHEL 9	`https://USERNAME:PASSWORD@archive.cloudera.com/p/cm7/patch/7.11.3.2-46642574/redhat9-ppc/yum`