

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 Release Notes

Date published: 2020-07-28

Date modified: 2026-03-31

The Cloudera logo is displayed in a bold, orange, sans-serif font. The word "CLOUDERA" is written in all caps, with a stylized 'E' that has a horizontal bar extending to the right.

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

Legal Notice

© Cloudera Inc. 2026. All rights reserved.

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

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

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

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

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

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

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

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

Contents

| | |
|--|-----------|
| Overview..... | 7 |
| What's new in Platform Support..... | 7 |
| What's New in Cloudera Runtime 7.3.2..... | 7 |
| Atlas..... | 8 |
| What's New in Cloud Connectors..... | 9 |
| Cruise Control..... | 9 |
| HBase..... | 10 |
| HDFS..... | 10 |
| Hive..... | 16 |
| Cloudera Data Explorer (Hue)..... | 17 |
| Iceberg..... | 18 |
| Impala..... | 20 |
| Kafka..... | 25 |
| Knox..... | 27 |
| Kudu..... | 28 |
| Livy..... | 28 |
| Navigator Encrypt..... | 29 |
| Oozie..... | 29 |
| Ozone..... | 29 |
| Phoenix..... | 33 |
| Ranger..... | 34 |
| What's New in Ranger KMS..... | 38 |
| Schema Registry..... | 39 |
| Solr..... | 39 |
| Spark..... | 39 |
| Spark Atlas Connector..... | 40 |
| Sqoop..... | 40 |
| Streams Messaging Manager..... | 40 |
| Streams Replication Manager..... | 40 |
| YARN and YARN Queue Manager..... | 41 |
| Zookeeper..... | 42 |
| Fixed Issues In Cloudera Runtime 7.3.2..... | 43 |
| Atlas..... | 43 |
| Avro..... | 46 |
| Cloud Connectors..... | 47 |
| Cruise Control..... | 47 |
| HDFS..... | 47 |
| HBase..... | 48 |
| Hive..... | 48 |
| Cloudera Data Explorer (Hue)..... | 56 |
| Impala..... | 57 |
| Fixed Issues in Apache Iceberg..... | 62 |
| Kafka..... | 65 |

| | |
|----------------------------------|----|
| Kudu..... | 65 |
| Knox..... | 67 |
| Livy..... | 68 |
| Navigator Encrypt..... | 68 |
| Oozie..... | 69 |
| Ozone..... | 70 |
| Parquet..... | 76 |
| Phoenix..... | 77 |
| Ranger..... | 77 |
| Ranger KMS..... | 81 |
| Schema Registry..... | 82 |
| Solr..... | 83 |
| Spark..... | 83 |
| Spark Atlas Connector..... | 84 |
| Sqoop..... | 84 |
| Streams Messaging Manager..... | 84 |
| Streams Replication Manager..... | 85 |
| Yarn and Yarn Queue Manager..... | 86 |
| ZooKeeper..... | 89 |

Known Issues In Cloudera Runtime 7.3.2..... 91

| | |
|---------------------------------------|-----|
| Atlas..... | 91 |
| Avro..... | 101 |
| Known Issues in Cloud Connectors..... | 101 |
| Cruise Control..... | 101 |
| HBase..... | 102 |
| HDFS..... | 102 |
| Hive..... | 102 |
| Cloudera Data Explorer (Hue)..... | 105 |
| Iceberg..... | 107 |
| Impala..... | 109 |
| Kafka..... | 115 |
| Knox..... | 118 |
| Kudu..... | 120 |
| Livy..... | 121 |
| Navigator Encrypt..... | 121 |
| Oozie..... | 121 |
| Ozone..... | 122 |
| Parquet..... | 126 |
| Phoenix..... | 126 |
| Ranger..... | 126 |
| Ranger KMS..... | 128 |
| Schema Registry..... | 128 |
| Solr..... | 129 |
| Spark..... | 129 |
| Spark Atlas Connector..... | 130 |
| Sqoop..... | 130 |
| Streams Messaging Manager..... | 131 |
| Streams Replication Manager..... | 133 |
| Yarn and Yarn Queue Manager..... | 134 |
| ZooKeeper..... | 136 |

Behavioral Changes in Cloudera Runtime 7.3.2..... 137

| | |
|----------------------------------|-----|
| Behavioral Changes in Atlas..... | 137 |
|----------------------------------|-----|

| | |
|--|------------|
| Behavioral Changes in Avro..... | 137 |
| Behavioral Changes in Cloud Connectors..... | 138 |
| Behavioral Changes in Cruise Control..... | 138 |
| Behavioral Changes in HBase..... | 138 |
| Behavioral Changes in HDFS..... | 141 |
| Behavioral Changes in Hive..... | 141 |
| Behavioral Changes in Cloudera Data Explorer (Hue)..... | 143 |
| Behavioral Changes in Impala..... | 143 |
| Behavioral Changes in Apache Iceberg..... | 145 |
| Behavioral Changes in Kafka..... | 146 |
| Behavioral Changes in Apache Knox..... | 147 |
| Behavioral Changes in Kudu..... | 147 |
| Behavioral Changes in Livy..... | 147 |
| Behavioral Changes in Oozie..... | 148 |
| Behavioral Changes in Ozone..... | 148 |
| Behavioral Changes in Parquet..... | 150 |
| Behavioral Changes in Apache Ranger..... | 150 |
| Behavioral Changes in Ranger KMS..... | 152 |
| Schema Registry..... | 152 |
| Behavioral Changes in Sqoop..... | 153 |
| Behavioral Changes in Streams Messaging Manager..... | 153 |
| Behavioral Changes in Streams Replication Manager..... | 154 |
| Behavioral Changes in Spark..... | 155 |
| Behavioral Changes in Spark Atlas Connector..... | 156 |
| Behavioral Changes in YARN..... | 156 |
| Behavioral Changes in ZooKeeper..... | 156 |
| | |
| Release Matrix..... | 157 |
| | |
| Cloudera Runtime Component Versions..... | 157 |
| Cloudera Runtime 7.3.2..... | 157 |
| | |
| Using the Cloudera Runtime Maven repository..... | 158 |
| Cloudera Runtime 7.3.2.0-957..... | 159 |
| | |
| Fixed Common Vulnerabilities and Exposures..... | 177 |
| Cloudera Runtime 7.3.2..... | 177 |
| | |
| API Modifications and Removals..... | 186 |
| Components with API differences in Cloudera Runtime 7.3.2..... | 186 |
| API Compatibility changes in 7.3.2 for Hadoop..... | 186 |
| API Compatibility changes in 7.3.2 for HBase..... | 188 |
| API Compatibility changes in 7.3.2 for Kafka..... | 201 |
| API Compatibility changes in 7.3.2 for Spark..... | 202 |
| API Compatibility changes in 7.3.2 for ZooKeeper..... | 202 |
| | |
| Deprecation Notices In Cloudera Runtime 7.3.2.0..... | 203 |
| Platform and OS..... | 204 |
| Deprecation Notices for Cruise Control..... | 204 |
| Deprecation Notices for Apache Ranger..... | 204 |

| | |
|--|-----|
| Deprecation Notices for Apache Kafka..... | 205 |
| Deprecation Notices for Streams Replication Manager..... | 205 |

Overview

This document provides a summary of the latest updates in Cloudera Runtime 7.3.2 and its Service Packs and Cumulative Hotfixes. It includes new features, improvements, known and fixed issues, technical previews, and more. For detailed, component-level information, see the [Cloudera documentation](#).

Understanding Cloudera platform versioning

Starting with Cloudera 7.3.2.0, versioning follows the <Base>.<Major>.<Minor>.<Patch> format:

- **Base:** Indicates the platform base version.
- **Major:** Represents the major version.
- **Minor:** Represents the minor version.
- **Patch:** Identifies a Cumulative Hotfix (CHF) or Service Pack (SP). A value of 0 in this position denotes a standard minor release.

For example, in version **7.3.2.0**, **7** is the base version, **3** is the major version, **2** is the minor version, and **0** indicates that it is an initial minor release (not a CHF or SP). For more details, see the [Cloudera Platform Support Policy](#).

What's new in Platform Support

You must be aware of the platform support changes for the Cloudera Runtime 7.3.2 release.

Cloudera Runtime 7.3.2

Platform Support Enhancements

- **New OS support:** Cloudera Runtime 7.3.2 now supports the following operating systems:
 - RHEL 9.6
 - Rocky Linux 9.6
 - SLES 15 SP6
 - Ubuntu 24.04

For more information about the operating system support, see [Cloudera Support Matrix](#).

- **New Database support:** Cloudera Runtime 7.3.2 now supports the following databases:
 - MariaDB 11.4
- **New JDK Version:** None

What's New in Cloudera Runtime 7.3.2

This version of Cloudera Runtime provides you with several new capabilities. Learn how the new features and improvements benefit you.



Note:

For information about Cloudera Base on premises 7.3.2 service packs upgrade paths, see [Supported in-place upgrade paths](#)

What's New in Apache Atlas

New features and functional updates for Atlas are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2 introduces new features of Atlas and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x [New Features](#).

Cloudera Runtime 7.3.2:

New React-Based User Interface for Apache Atlas

Apache Atlas now features a redesigned React-based user interface that offers enhanced usability and streamlined metadata management. You can switch between the Classic and New UI experiences. The new interface introduces an improved search panel that automatically lists all available entity types, classifications, and glossary terms, with one-click access to relevant members. Enhanced filtering capabilities allow users to show empty service types, unused classifications, and toggle between category or term views in the glossary. Additionally, entities and classifications can now be displayed in a collapsed flat tree view for simplified navigation of complex metadata hierarchies. For more information, see [Apache Atlas dashboard tour](#).

Apache Atlas component upgraded to 2.4.0

The Atlas runtime component is upgraded from 2.1.0 to 2.4.0. Several stability and correctness fixes are included from the upstream release for bugs, including user interface improvements for classification propagation settings.

Atlas Auto-Purging introduced



Important: This feature is available as technical preview and is under entitlement. To obtain the required entitlement, contact your Cloudera Account Representative.

The automated entity auto-purging feature addresses potential performance and storage issues caused by the previous manual purge strategy. Atlas preserves metadata by only marking entities as deleted. This leads to query performance degradation and increased storage usage as soft-deleted entities accumulate. The soft-deleted items could be only manually deleted by using the [PUT /admin/purge/ API](#) call, however, this API call leaves behind the column lineage entities for soft-deleted process entities.

The new, cron-based system can be configured to clean up obsolete process entities, including their column lineage entities, that are no longer relevant. This prevents sparse graphs and significantly improves metadata hygiene and query performance. For more information, see [Atlas Auto-Purging overview](#).

Support replication of Atlas data from on-prem to on-prem



Important: This feature is available as technical preview and is under entitlement. To obtain the required entitlement, contact your Cloudera Account Representative.

Atlas now supports asynchronous import of metadata using Kafka. Previously, the only import mechanism was synchronous: the HTTP connection remained open until the entire import completed, causing timeouts for large datasets and making concurrent imports fragile. With asynchronous import, a client submits an import request that is immediately staged and queued as a Kafka message, and receives an import ID in response without waiting for processing to finish. Atlas processes the import in the background and persists the request state, including received time, processing start time, completion time, and outcome.

The following new REST API endpoints are available:

- [POST /api/atlas/admin/async/import](#) — submit an asynchronous import; returns immediately with an import ID
- [GET /api/atlas/admin/async/import/status](#) — list all async import statuses

- GET `/api/atlas/admin/async/import/status/{importId}` — get the status of a specific import
- DELETE `/api/atlas/admin/async/import/{importId}` — abort a specific queued import

Atlas upgraded to use JDK 17

Atlas now runs on Java 17, upgraded from Java 8. JDK 17 is a Long-Term Support (LTS) release that brings improved performance, enhanced security, and better long-term maintainability. Key benefits for Atlas users include:

- Improved garbage collection, resulting in lower latency and more efficient memory usage for metadata-intensive workloads.
- Stronger cryptographic algorithms reducing security vulnerabilities.
- Long-term support guaranteed until at least 2029, ensuring continued security patches.

Logback introduced as logging framework

Apache Atlas now uses Logback as its logging framework, replacing Log4j2. This change enhances security and simplifies log management. It also enables the user to add any new properties overriding existing properties.

- **Simplified Configuration:** Streamlined logging setup is introduced with native XML configuration instead of the `.properties` file.
 - Go to [Cloudera Manager Atlas Server XML Override](#) to replace the complete configuration file.
- Configuration still remains the same, as file size and rotation.

What's New in Cloud Connectors

Learn about the new features of Cloud Connectors in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces new features of Cloud Connectors and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.700. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see [New Features](#).

Rebase to Hadoop 3.4.2

This release of the Cloud Connectors is based on Hadoop 3.4.2. See the following upstream resources for more information on the changes:

- [Apache Hadoop 3.4.2](#)
- [Apache Hadoop 3.4.1](#)
- [Apache Hadoop 3.4.0](#)

What's New in Cruise Control

New features and functional updates for Cruise Control are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

New configuration parameter for controlling IP stack preference

A new `cc.additional.java.options` configuration parameter is available on the Cruise Control configuration page in Cloudera Manager. The default value sets the IP protocol to IPv4.

New `intra.broker.goals` configuration for Cruise Control

Cloudera Manager introduces a new `intra.broker.goals` configuration for Cruise Control. The default value includes `com.linkedin.kafka.cruisecontrol.analyzer.goals.IntraBrokerDiskCapacityGoal` and `com.linkedin.kafka.cruisecontrol.analyzer.goals.IntraBrokerDiskUsageDistributionGoal`.

This has an effect on the existing Default Goals (`default.goals`) configuration, which must be a subset of Supported Goals and Supported Intra Broker Goals.

Additionally, the `intra.broker.goals` configuration no longer needs to be defined in an advanced configuration snippet if done previously.

What's New in Apache HBase

New features and functional updates for HBase are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2:

IPv6 support for HBase

Starting with the 7.3.2 release, HBase supports IPv6 with dual-stack functionality, allowing seamless communication over both IPv4 and IPv6 networks. This capability improves network scalability, future-proofs deployments, and enhances overall platform security.

For more information, see [Enabling IPv6 support for HBase](#).

HBase JDK 17 upgrade

HBase supports JDK 17 starting with the 7.3.2 release. Upgrade to JDK 17 to gain improved performance, increased security, modern language features, and long-term support, ensuring your application remains competitive and maintainable.

Apache HBase component upgraded to 2.6.3

The HBase runtime component is upgraded from 2.4.17 to 2.6.3. This release incorporates stability and correctness fixes from the upstream, which deliver performance and maintenance improvements, and enhanced client connectivity support.

HBase does not support file-based implementation, but the project plans to support the table-based implementation after rebasing to the HBase 3.x version.

What's New in HDFS

Learn about the new features of HDFS in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Support for G1 Garbage Collector (G1GC) with JDK 17

When running HDFS on JDK 17, all HDFS server processes now automatically use the G1 Garbage Collector (G1GC). No action is required for most deployments.

Hadoop rebase summary

In Cloudera Runtime 7.3.2, Apache Hadoop is rebased to version 3.4.1. The Apache Hadoop upgrade improves overall performance and includes all the new features, improvements, and bug fixes from versions 3.2, 3.3, and 3.4.

Table 1: New features added from Apache Hadoop 3.2 to 3.4 versions

| Apache Hadoop version | Apache Jira | Name | Description |
|-----------------------|------------------------------|--|---|
| 3.2 | HDFS-10285 | Storage Policy Satisfier in HDFS | <p>StoragePolicySatisfier (SPS) allows you to track and fulfill the storage policy requirement of a given file or directory in HDFS. You can specify a file or directory path for SPS to evaluate by running the <code>hdfs storagepolicies -satisfyStoragePolicy -path <path></code> command or invoking the <code>HdfsAdmin#satisfyStoragePolicy(path)</code> API.</p> <p>If blocks have storage policy mismatches, SPS moves the replicas to a different storage type to fulfill the storage policy requirements.</p> <p>Because API calls go to the NameNode to track the invoked satisfier path (iNodes), you must enable SPS on the NameNode by setting the <code>dfs.storage.policy.satisfier.mode</code> property to <code>external</code> in the <code>hdfs-site.xml</code> file. You can be disable the configuration dynamically without restarting the NameNode.</p> <p>SPS must be started outside the NameNode using the <code>hdfs -daemon start sps</code> command. To run the Mover tool explicitly, you must disable the SPS first. For more information, see the <i>Storage Policy Satisfier (SPS)</i> section in the Archival Storage .</p> <p>By default, Storage Policy Satisfier is disabled in Cloudera HDFS.</p> |
| 3.3 | HDFS-14118 | DNS resolution for nameservices to IP addresses | HDFS clients can use a single domain name to discover servers such as namenodes, routers, or observers instead of explicitly listing out all hosts in the configurations. |
| 3.3 | HDFS-13783 | Long-running service process for the Balancer | <p>Adds the <code>-asService</code> parameter to the Balancer CLI. This parameter enables the Balancer to run as a long-running service process for easier monitoring.</p> <p>By default, this service mode is disabled in Cloudera HDFS.</p> |
| 3.3 | HADOOP-16398 | Hadoop metrics export to Prometheus | <p>If the <code>hadoop.prometheus.endpoint.enabled</code> configuration is set to true, Prometheus-friendly formatted metrics can be obtained from the <code>/prom</code> endpoint of Hadoop daemons. The default value is false.</p> <p>By default, the configuration is disabled in Cloudera HDFS.</p> |
| 3.3 | HDFS-13571 | Dead node detection | When an unresponsive node blocks a <code>DFSInputStream</code> , the system detects the failure and shares this information to other streams in the same <code>DFSClient</code> . This prevents the remaining streams from attempting to read from and be blocked by the dead node . |
| 3.4 | HADOOP-17010 | Queue capacity weights in FairCall Queue | <p>When the <code>FairCallQueue</code> feature is enabled, you can specify capacity allocation among all sub-queues using the <code>ipc.<port>.callqueue.capacity.weights</code> configuration. The value of this configuration is a comma-separated list of positive integers, each of which specifies the weight associated with the sub-queue at that index. The number of weights in this list must match the IPC scheduler priority levels defined in <code>scheduler.priority.levels</code>.</p> <p>By default, each sub-queue is associated with a weight of 1, meaning all sub-queues are allocated with the same capacity.</p> <p>This configuration is optional in Cloudera HDFS.</p> |
| 3.4 | HDFS-13183 | Standby NameNode processing for the getBlocks requests to reduce the active load | <p>Enables the balancer to redirect <code>getBlocks</code> request to a standby NameNode, thus reducing the performance impact of the balancer to the active NameNode.</p> <p>The feature is disabled by default. To enable it, set the <code>dfs.ha.allow.stale.reads</code> configuration of the balancer to true in the <code>hdfs-site.xml</code> file.</p> |

| Apache Hadoop version | Apache Jira | Name | Description |
|-----------------------|----------------------------|---------------------------------------|--|
| 3.4 | HDFS-15025 | NVDIMM storage media support for HDFS | Adds the NVDIMM storage type and the ALL_NVDIMM storage policy for HDFS. The NVDIMM storage type is for non-volatile random-access memory storage media whose data survives when the DataNode restarts. |
| 3.4 | HDFS-15098 | SM4 encryption method for HDFS | The SM4/CTR/NoPadding encryption codec is now added. It requires OpenSSL 1.1.1. or higher version for native implementation. |
| 3.4 | HDFS-15747 | RBF: Renaming across sub-namespaces | Renames multiple namespaces on the Federation balance tool. |
| 3.4 | HDFS-15547 | Dynamic disk-level tiering | Archival Storage allows the configuration of DISK and ARCHIVE types on the same device. It enables optimizing disk I/O in heterogeneous environments. |
| 3.4 | HDFS-16595 | Slow peer metrics | NameNode metrics that represent slownode JSON now include three important factors, namely median, median absolute deviation, and upper latency limit, which can help you determine how urgently a given slownode requires manual intervention. |

Table 2: Improvements added from Apache Hadoop 3.2 to 3.4 versions

| Apache Hadoop version | Apache Jira | Name | Description |
|-----------------------|------------------------------|---|---|
| 3.3 | HDFS-7133 | Clearance of the namespace quota on / | Namespace quota on root / can be cleared now. |
| 3.3 | HADOOP-13363 | Upgrade of protobuf from 2.5.0 to the highest version | Upgrade protobuf from 2.5.0 to the highest version. |
| 3.3 | HADOOP-16670 | Removal of Submarine code from Hadoop codebase | The Submarine subproject has been moved to its own Apache Top Level Project. Therefore, the Submarine code has been removed from Hadoop 3.3.0 and higher versions. For more information, see https://submarine.apache.org/ . |
| 3.3 | HADOOP-17021 | Addition of the -concat fs command | The <code>hadoop fs</code> utility includes a <code>-concat</code> command available on all filesystems that support the <code>concat</code> API, including HDFS and WebHDFS. |
| 3.4 | HADOOP-17222 | Creation of socket address leveraging the URI cache | <p>The DFS client can use the newly added URI cache when creating a socket address for read operations. When enabled, creating a socket address will use cached URI object based on host:port to reduce the frequency of URI object creation.</p> <p>To enable it, set the following configuration key to true:</p> <pre><property> <name>dfs.client.read.uri.cache.enabled</name> <value>true</value> </property></pre> <p>By default, the configuration is disabled in Cloudera HDFS.</p> |
| 3.4 | HDFS-15814 | Configurable parameters for DataNodeDiskMetrics | Makes certain parameters configurable for DataNode DiskMetrics. |

| Apache Hadoop version | Apache Jira | Name | Description |
|-----------------------|--|--|--|
| 3.4 | HADOOP-16747 | Support for Python 3 | Supports Python 3. |
| 3.4 | HADOOP-17514 | Removal of the <code>trace</code> subcommand from the Hadoop CLI | The <code>trace</code> subcommand of the Hadoop CLI has been removed as a follow-up of removal of the TraceAdmin protocol. |
| 3.4 | HADOOP-16524 | Automatic keystore reloading for HttpServer2 | <p>Adds auto-reload of the keystore.</p> <p>Adds the following new configuration (default is 10 seconds):</p> <pre>ssl.{0}.stores.reload.interval</pre> <p>The refresh interval is used to check if either the truststore or keystore certificate file has changed.</p> |
| 3.4 | HDFS-15975 | Replacement of LongAdder with AtomicLong | This update changes public fields in DFSHedgedReadMetrics. If you are using the public member variables of DFSHedgedReadMetrics, you must access them using the public API. |
| 3.4 | HADOOP-17956 | Replacement of all default Charset usage with UTF-8 | All of the default charset usages have been replaced to UTF-8. If the default charset of your environment is not UTF-8, the behavior can be different. |
| 3.4 | HDFS-16278 , HDFS-16285 , and HDFS-16419 | Cross-platform support for HDFS snapshot tools | Supports HDFS snapshot tools cross-platform. |
| 3.4 | HDFS-15382 | Division of the single FsDatasetImpl lock into volume-grained locks | Throughput is one of the core performance evaluations for the DataNode instance. However, it does not reach the best performance, especially for Federation deployments, all the time because of the global coarse-grain lock. HDFS-16534 , HDFS-16511 , HDFS-15382 and HDFS-16429 aim to split the global coarse-grain lock into a fine-grain lock, which is a double-level lock for blockpool and volume, to improve the throughput and avoid lock impacts between blockpools and volumes. |
| 3.4 | HADOOP-18188 | Support for the <code>touch</code> command for directories | Previously, <code>hadoop fs -touch</code> command threw a PathIsDirectoryException error for a directory. Now, the command supports directories and will not throw that exception. |
| 3.4 | HADOOP-13332 | Removal of Jackson 1.9.13 and migration to the 2.x codebase | Remove Jackson 1.9.13 and switch all Jackson code to the 2.x codebase. |
| 3.4 | HADOOP-18332 | Removal of the <code>rs-api</code> dependency by downgrading Jackson to 2.12.7 | Downgrades Jackson from 2.13.2 to 2.12.7 to fix class conflicts in downstream projects. This version of Jackson does contain the fix for CVE-2020-36518. |
| 3.4 | HADOOP-18442 | Removal of the <code>hadoop-openstack</code> module | The <code>swift://</code> connector for Openstack support has been removed because of fundamental limitations. such as Swift's handling of files greater than 4 GB. Because most object store services now export a subset of the S3 protocol, use the <code>s3a</code> connector instead. The <code>hadoop-openstack</code> JAR remains as an empty artifact to prevent build failures for projects that still declare the JAR as a dependency. |
| 3.4 | HADOOP-16206 | Migration from commons-logging to reload4j | Migrates commons-logging to reload4j. |

| Apache Hadoop version | Apache Jira | Name | Description |
|-----------------------|------------------------------|---|--|
| 3.4 | HADOOP-17524 | Removal of EventCounter and Log counters from JVMetrics | Removes the log counter from JVMetrics because this feature strongly depends on the Log4J 1.x API. |
| 3.4 | HADOOP-18206 | Cleanup of commons-logging references in the codebase | Cleans up the commons-logging references in the codebase. |
| 3.4 | HADOOP-18820 | AWS SDK v2: Optional v1 bridging support | The v1 aws-sdk-bundle JAR is removed. It is only required for third party applications or when using v1 SDK AWSCredentialsProvider classes. Standard providers automatically migrate from the v1 to v2 classes. Therefore this change only affects third-party providers or environments using very esoteric classes in the V1 SDK . For more information, see Upgrading S3A to AWS SDK V2 . |
| 3.4 | HADOOP-18876 | ABFS: Default buffer change from disk to ByteBuffer for fs.azure.data.blocks.buffer | The default value for the fs.azure.data.blocks.buffer configuration is now changed from disk to ByteBuffer. This speeds up writing to Azure storage, at the risk of running out of memory, especially if many threads are writing to ABFS at the same time and the upload bandwidth is limited. If jobs do run out of memory while writing to ABFS, change the option back to disk. |
| 3.4 | HADOOP-18993 | S3A: Addition of the fs.s3a.classloader.isolated option | To load custom implementations of AWS Credential Providers through user-provided jars, set the <code>{{fs.s3a.extensions.isolated.classloader}}</code> configuration to <code>{{false}}</code> . |
| 3.4 | HADOOP-19120 | ABFS: Adaptation of ApacheHttpClient as a network library | <p>Apache HttpClient 4.5.x is a new implementation of HTTP connections, supporting a large, configurable pool of connections along with the ability to limit their lifespan.</p> <p>Configure the networking library using the fs.azure.networking.library configuration option.</p> <p>The following are the supported values:</p> <ul style="list-style-type: none"> JDK_HTTP_URL_CONNECTION Uses the JDK networking library [Default] APACHE_HTTP_CLIENT Uses Apache HttpClient <p> Important: If you switch the networking library to the Apache HttpClient, you must ensure that Apache HttpClient and HttpClient are on the classpath.</p> |

Table 3: Issues fixed between Apache Hadoop versions 3.2 to 3.4

| Apache Hadoop version | Apache Jira | Name | Description |
|-----------------------|------------------------------|--|--|
| 3.3 | HDFS-14396 | Image loading failure from FSImageFile during downgrades from 3.x to 2.x | During a rolling upgrade from Hadoop 2.x to 3.x, NameNode could not persist erasure coding information, and therefore a user could not start using the erasure coding feature until the finalize operation was complete. |
| 3.3 | HADOOP-15355 | SFTPConnectionPool connection leakage | SFTPConnectionPool connection leakage occurred. |

| Apache Hadoop version | Apache Jira | Name | Description |
|-----------------------|------------------------------|--|--|
| 3.3 | HDFS-14845 | Bypass of AuthenticationFilterInit for HttpFSWebServer | HttpFS previously ignored standard Hadoop authentication settings. HttpFS now honors <code>hadoop.http.authentication.*</code> configurations and the previous <code>httpfs.authentication.*</code> configurations are deprecated. If both configurations are set, the deprecated <code>httpfs.authentication.*</code> configurations take precedence for backward compatibility. |
| 3.4 | HDFS-15380 | RBF: Fetch failure for the real remote IP in RouterWebHdfsMethods | RBF could not fetch the real remote IP in RouterWebHdfsMethods. |
| 3.4 | HDFS-15281 | ZKFC host binding fallback from the <code>dfs.namenode.rpc-bind-host</code> to <code>dfs.namenode.rpc-address</code> | ZKFC previously ignored the <code>dfs.namenode.rpc-bind-host</code> configuration. ZKFC now correctly binds the host address using <code>dfs.namenode.servicerpc-bind-host</code> if configured, falling back to <code>dfs.namenode.rpc-bind-host</code> . If neither configuration exists, ZKFC binds to the NameNode RPC server address (acting as <code>dfs.namenode.rpc-address</code> configuration). |
| 3.4 | HADOOP-18237 | Apache Xerces Java upgrade to 2.12.2 | An older version of Apache Xerces Java contained vulnerabilities. Apache Xerces has been updated to 2.12.2 to resolve CVE-2022-23437. |
| 3.4 | HADOOP-18629 | CryptoOutputStream::close leakage | A resource leak occurred in <code>CryptoOutputStream::close</code> when encountering encrypted zones combined with quota exceptions occur. |
| 3.4 | HADOOP-18329 | Missing support for Semeru OE JRE 11.0.15.0 | Support was missing for Semeru Runtimes, in which, due to vendor name-based logic and changes in the <code>java.vendor</code> property, failures could occur on Java 11 Runtimes 11.0.15.0 and higher versions. |
| 3.4 | HADOOP-19225 | S3A: Recovery failure for S3A multipart block upload attempts Status Code: 400; Error Code: RequestTimeout | S3A upload operations failed to recover when the store returned a 500 error. These operations now successfully recover. You can control this retry behavior for 50x errors, excluding 503 throttling events, which remain independently processed, using the <code>fs.s3a.retry.http.5xx.errors</code> property. The default value is true. |

Hadoop interface standards and API changes

Hadoop provides many [compatibility](#) promises, including the following guarantees:

- Java APIs marked [Stable](#) must not change within a major release.
- Hadoop wire protocols are classified as [Private](#) and [Stable](#), and remain forward-compatible across minor releases within a specific major version.
- The exposed Hadoop REST APIs are generally classified as [Public](#) and [Evolving](#). However, within a specific API version number, they are treated as Public and Stable, meaning no incompatible changes are permitted.

The following library changes affect Cloudera components:

- The commons-lang (2) has been migrated to commons-lang3, which uses a new package name. Dependent Cloudera components must update their imports.
- Guava and some instances of Protobuf now come as shaded JARs from [hadoop-thirdparty](#).
- A new bouncycastle version, including Java 17 and 21 class variants, is now available. Tools such as Gradle and maven-shade-plugin must be updated to higher versions or excluded when shading.

- Switched from `javax.activation` to `jakarta.activation`.
- `org.apache.kerby.kerb-simplekdc` is no longer imported by Hadoop, so components that depend on it must explicitly include it.
- Migrated from `org.codehaus.jackson (1.x)` to `com.fasterxml.jackson (2.x)`.
- The Hadoop `EventCounter` class was removed in [HADOOP-17524](#) with no replacement because it relied on deprecated `log4j 1.x` functionality not present in 2.x.
- Removed `commons-logging`. Hadoop now exclusively configures and uses `reload4j` now.

The following Java API changes affect Cloudera components:

- [HADOOP-16645](#) – Switches S3A APIs to return a `StoreContext` rather than provide direct access to `S3AFileSystem`.
- [HADOOP-18206](#) – Removed the `IOUtils#cleanup` method from Public-Evolving API as part of the migration away from `commons-logging`.
- [HDFS-16982](#) – Renamed `MutableQuantiles.quantiles` to `MutableQuantiles.QUANTILES`.

What's New in Apache Hive

New features and functional updates for Hive are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces new features of Hive and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see [New Features](#).

Hive user experience enhancements

Cloudera now provides several Hive user experience enhancements:

- Improved error handling for the `SHOW PARTITIONS` command: The `SHOW PARTITIONS` command now returns a concise execution error instead of a full stack trace when you run it against a non-partitioned table. This change simplifies the output and helps you identify configuration issues more quickly. (HIVE-26926)
- Enhanced error messages for the `STORED BY` clause: The error message displayed when you provide an invalid identifier or literal in the `STORED BY` clause is now more informative. This improvement provides better guidance if you mistakenly use `STORED BY` instead of `STORED AS` or provide an incorrect storage format. (HIVE-27957)
- Enhanced task attempt log clarity: The task attempt log now includes explicit context for boolean values to improve readability. You can now clearly identify whether a task attempt is guaranteed or not within the log messages. (HIVE-28246)
- Enabled vectorized mode support for custom UDFs: You can now use custom User-Defined Functions (UDFs) in vectorized execution mode. This enhancement ensures that custom functions, such as `TIME_PARSE`, integrate correctly with vectorized query plans to improve processing performance. (HIVE-28830)
- Resolved `NullPointerException` in `TezSessionPoolManager`: A `NullPointerException` (NPE) that occurred in the `TezSessionPoolManager` when the resource plan was null is now resolved. This fix improves the stability of the Tez session pool when updating triggers from an active resource plan. (HIVE-29007)

Apache Jira: [HIVE-26926](#), [HIVE-27957](#), [HIVE-28246](#), [HIVE-28830](#), [HIVE-29007](#)

Dropping Hive Metastore statistics

You can now drop statistics associated with tables, partitions, and columns in the Hive Metastore (HMS). This feature is particularly useful during migration or replication processes where large volumes of statistical data—generated for every table, partition, and column combination—can

significantly increase copy times. By removing unnecessary statistics, you can improve operational efficiency and reduce data transfer overhead.

Apache Jira: [HIVE-28655](#)

Enhanced Hive Metastore notification fetching with table filters

The Hive Metastore (HMS) notification fetch API now supports optional database and table name filters. This enhancement allows external engines and clients to fetch pending events specifically for a given table or list of tables, rather than scanning all notifications. Additionally, an index on the table name is now included in the HMS notification log to optimize query performance and prevent table scans. This change enables more efficient metadata synchronization and improves performance for queries involving multiple tables.

Apache Jira: [HIVE-27499](#)

New command to display HiveServer2 and Hive Metastore connections

You can now use the `SHOW PROCESSLIST` command to display active operations and connection details for HiveServer2 (HS2) and Hive Metastore (HMS) instances. This feature provides a view of current sessions, including user names, IP addresses, query IDs, and execution states, similar to the process list functionality in MySQL. This command helps you troubleshoot stuck queries, monitor service load, and identify inappropriate connections for termination.

Apache Jira: [HIVE-27829](#)

Upgrading Calcite

Hive has been upgraded to Calcite version 1.33. This upgrade introduces various query optimizations that can improve query performance.

Hive on ARM Architecture

Hive is now fully supported on ARM architecture instances, including AWS Graviton and Azure ARM. This enables you to run your Hive workloads on more cost-effective and energy-efficient hardware.

ZooKeeper SASL authentication for Hive clients

You can now configure Hive clients to authenticate with a ZooKeeper ensemble that enforces Simple Authentication and Security Layer (SASL). By using the new `hive.zookeeper.client.sasl.enforce` property, JDBC and HS2 clients can successfully establish sessions in Kerberized environments that require secure service discovery and locking. For more information, see [Configuring Zookeeper SASL for Hive](#)

Impala now supports OAuth Authentication

Cloudera now provides support for OAuth authentication using OAuth JWT bearer tokens. For more information, see:

[Impala OAuth Authentication](#)

What's New in Cloudera Data Explorer (Hue)

New features and functional updates for Cloudera Data Explorer (Hue) are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces new features of Cloudera Data Explorer (Hue) and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see [New Features](#).

Product Branding Update

The product component previously known as **Hue** is now renamed to Cloudera Data Explorer (Hue). This change reflects UI and an updated branding initiative and will be rolled out in phases.

As part of this release, you can notice the following changes:

- A new logo in the UI
- The service name updated to Data Explorer in the UI
- The new product name reflected in documentation

Some UI references might still display the previous name as the branding update is completed incrementally in future releases.

No functional impact is associated with this change. All existing configurations, workflows, and integrations continue to work as before.

Data Explorer JDK 17 upgrade

Data Explorer supports JDK 17. Upgrading to JDK 17 ensures your application stays competitive and maintainable by providing improved performance, increased security, modern language features, and long-term support.

Data Explorer IPv6 support

Data Explorer supports IPv6 with dual-stack functionality, allowing seamless communication over both IPv4 and IPv6 networks. This capability improves network scalability and enhances overall platform security.

For more information, see [Enabling Dual-Stack support for Data Explorer](#).

Python 3.11 support in Data Explorer

Data Explorer now supports Python 3.11 as the only supported version across all operating systems. Python 3.8, 3.9, and 3.10 are no longer supported.

Enhanced session security for Data Explorer

Data Explorer now includes security for the session ID (sessionid) cookie. This enhancement helps prevent unauthorized access, resulting in data exposure, unauthorized query execution, and job submission across connected Data Explorer services.

For more information, see [Securing Data Explorer sessions](#).

Improved navigation pane interpreter visibility

Previously, enabling an application such as SparkSQL automatically displayed every associated default interpreter, which often cluttered the UI. Data Explorer now prioritizes a streamlined interface by no longer displaying all available interpreters by default. If your workflow requires the previous behavior, you can revert to displaying all available interpreters by enabling the `enable_all_interpreters` configuration flag.

For more information, see [Enabling multiple editors in Cloudera Data Explorer \(Hue\)](#).

Global Hive JDBC URL for Oozie workflows

Data Explorer now supports a global configuration for Hive JDBC connection strings used in Oozie workflow. Data Explorer derived these URLs from the default Hive configuration, requiring manual, per-workflow overrides for specialized setups such as High Availability (HA) quorums. You can now use the `oozie_hs2_jdbc_url` configuration property to set a single, cluster-wide JDBC URL.

For more information, see [Configuring global Hive JDBC URL for Oozie](#).

What's New in Apache Iceberg

Learn about the new features of Iceberg in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces new features of Iceberg and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see [New Features](#).

Cloudera Lakehouse Optimizer for Iceberg table optimization

In Cloudera Runtime 7.3.2 and higher versions, you can use Cloudera Lakehouse Optimizer service in Cloudera Manager to automate the Iceberg table maintenance tasks.

Cloudera Lakehouse Optimizer provides automated Iceberg table maintenance, through Spark jobs, for Iceberg tables in Cloudera Open Data Lakehouse. It simplifies table management, improves query performance, and reduces operational costs.

You can add the service to an existing Cloudera Base on premises 7.3.2 or higher versions cluster in Cloudera Manager 7.13.2 or higher versions, or you can create a dedicated cluster and then add the service. You must ensure that the cluster contains all the required services. After you finish configuring the service, you can use the Cloudera Lakehouse Optimizer service REST APIs to define the Cloudera Lakehouse Optimizer policies, perform policy management, and run other Iceberg table optimization operations.

For more information, see [Cloudera Lakehouse Optimizer](#).

Integrate Iceberg scan metrics into Impala query profiles

Iceberg scan metrics are now integrated into the Frontend section of Impala query profiles, providing deeper insight into query planning performance for Iceberg tables.

The query profile now displays scan metrics from Iceberg's `planFiles()` API, including total planning time, counts of data/delete files and manifests, and the number of skipped files.

Metrics are displayed on a per-table basis. If a query scans multiple Iceberg tables, a separate metrics section will appear in the profile for each one.

Apache Jira: [IMPALA-13628](#)

Delete orphan files for Iceberg tables

You can now use the following syntax to remove orphan files for Iceberg tables:

```
-- Remove orphan files older than '2022-01-04 10:00:00'.
ALTER TABLE ice_tbl EXECUTE remove_orphan_files('2022-01-04 10:
00:00');

-- Remove orphan files older than 5 days from now.
ALTER TABLE ice_tbl EXECUTE remove_orphan_files(now() - interval
5 days);
```

This feature removes all files from a table's data directory that are not linked from metadata files and that are older than the value of `older_than` parameter. Deleting orphan files from time to time is recommended to keep the size of a table's data directory under control.

Apache Jira: [IMPALA-14492](#)

Allow forced predicate pushdown to Iceberg

Since [IMPALA-11591](#), Impala has optimized query planning by avoiding predicate pushdown to Iceberg unless it is strictly necessary. While this default behavior makes planning faster, it can miss opportunities to prune files early based on Iceberg's file-level statistics.

A new table property, `impala.iceberg.push_down_hint` is introduced, which allows you to force predicate pushdown for specific columns. The property accepts a comma-separated list of column names, for example, `'col_a, col_b'`.

If a query contains a predicate on any column listed in this property, Impala will push that predicate down to Iceberg for evaluation during the planning phase.

Apache Jira: [IMPALA-14123](#)

UPDATE operations now skip rows that already have the desired value

The UPDATE statement for Iceberg and Kudu tables is optimized to reduce unnecessary writes.

Previously, an UPDATE operation would modify all rows matching the WHERE clause, even if those rows already contained the new value. For Iceberg tables, this resulted in writing unnecessary new data and delete records.

With this enhancement, Impala automatically adds an extra predicate to the UPDATE statement to exclude rows that already match the target value.

Apache Jira: [IMPALA-12588](#)

What's New in Apache Impala

New features and functional updates for Impala are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces new features of Impala and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see [New Features](#).

Hierarchical metastore event processing (Preview)

Impala now supports a multi-layered, hierarchical approach to metastore event processing to improve synchronization speed and handle event dependencies more efficiently. By enabling this feature, you can segregate events based on their dependencies and process them independently through a system of database and table event executors. This method reduces synchronization time for Hive Metastore (HMS) events by allowing parallel processing while maintaining linearizability for specific tables.

For more information, see [Hierarchical metastore event processing](#)

Impala AES encryption and decryption support

Impala now supports AES (Advanced Encryption Standard) encryption and decryption to work better with other systems. AES-GCM is the default mode for strong security, but you can also use other modes like CTR, CFB, and ECB for different needs. This feature works with both 128-bit and 256-bit keys and includes checks to keep your data safe and confidential.

For more information, see [AES encryption and decryption support](#)

Apache Jira: [IMPALA-13039](#)

Ubuntu 24.04 support

You can now build and run Impala on Ubuntu 24.04.

Dual-stack networking support

You can now configure Impala to support dual-stack networking, allowing the service to handle both IPv4 and IPv6 traffic simultaneously. This update includes new configuration properties in Cloudera Manager and support for dual-stack load balancing with HAProxy.

For more information, see [Impala dual stack IPv6](#) and [Impala dual stack HA Proxy](#)

Default support for Java 17

Cloudera now provides Impala using Java 17 for builds and runtime environments.

OpenTelemetry integration for Impala

Cloudera now provides OpenTelemetry (OTel) support to help you see query performance and troubleshoot issues. This new feature, collects and exports query telemetry data as OpenTelemetry traces to a central OpenTelemetry compatible collector. The integration is designed to have a minimal impact on performance because it uses data already being collected and handles the export in a separate process. For more information, see [OpenTelemetry support for Impala](#)

Apache Jira: [IMPALA-13234](#)

Caching intermediate query results

Cloudera now supports caching intermediate results to improve query performance and resource efficiency for repetitive workloads. By storing results at various locations within the SQL plan tree, the system can reuse computation for similar queries even when they are not identical, provided the underlying data and settings remain unchanged. For more information, see [Caching intermediate results](#)

Impala-shell and Impyla now supports Python 3.12

Cloudera now provides support for Python 3.12 in `impala-shell` and `Impyla`.

Apache Jira: [IMPALA-14452](#)

Query cancellation supported during analysis and planning

This new feature allows you to cancel Impala queries even while they are in the Frontend stage, which includes analysis and planning. Previously, you could not cancel a query while it was waiting for operations like loading metadata from the Catalog Server. With this update, Impala now registers the planning process and can interrupt it to cancel the query.

Apache Jira: [IMPALA-915](#)

Improved memory estimation and control for large queries

Impala now uses a more realistic approach to memory estimation for large operations like SORT, AGGREGATION, and HASH JOIN.

Previously, these operations could severely overestimate their memory needs (sometimes requesting terabytes) when row counts were misestimated. This often caused the Admission Controller to reject the query outright, even though the operation could easily handle the data by writing (spilling) to disk.

The system now considers the operator's ability to spill data to disk and caps the memory estimate based on your cluster's actual memory limits (like `MEM_LIMIT`). This change makes more large queries admissible and allows them to run successfully.

New Control Option: A new query option, `MEM_ESTIMATE_SCALE_FOR_SPILLING_OPERATOR` (a scale from 0.0 (exclusive) to 1.0 (inclusive)), is introduced, giving you control over this behavior:

- Higher values (closer to 1.0): Tells Impala to reserve more memory, increasing the chance of faster, in-memory execution.
- Lower values (closer to 0.0, but NOT 0.0): Tells Impala to request less memory, increasing the chance your query is admitted, but potentially leading to more spilling to disk (slower execution).
- The default value is 0.0 which disables this feature, and revert Impala planner to old behavior.

Apache Jira: [IMPALA-13333](#)

Expose query cancellation status to UDF interface

Impala now exposes the query cancellation status to the User-Defined Function (UDF) interface. This new feature allows complex or time-consuming UDFs to periodically check if the query has been cancelled by the user. If cancellation is detected, the UDF can stop its work and fail fast.

This significantly reduces the time you have to wait to stop a long-running query that is stuck inside a UDF evaluation.

Apache Jira: [IMPALA-13566](#)

Expanded compression levels for ZSTD, and ZLIB

Impala has extended the configurable range of compression levels for ZSTD, and ZLIB (GZIP/DEFLATE) codecs. This enhancement allows for better optimization of the trade-off between compression ratio and write throughput.

- ZSTD: Supports a wider range, including negative levels, up to 20.
- ZLIB (GZIP, DEFLATE): Supports levels from 1 (default) to 9 (best compression).

These levels are applied via the `compression_codec` query option.

JIRA Issue: [IMPALA-13923](#)

Constant folding is now supported for non-ASCII and binary strings

Previously, the query planner could not apply the optimization known as constant folding if the resulting value contained non-ASCII characters or was a non-UTF8 binary string. This failure meant that important query filters could not be simplified, which prevented key performance optimizations like predicate pushdown to the storage engine (e.g., Iceberg or Parquet stat filtering).

The planner is updated to correctly handle and fold expressions resulting in valid UTF-8 strings (including international characters) and binary byte arrays. This allows Impala to push down more filters, significantly improving the performance of queries that use non-ASCII string literals or binary data in their filters.

JIRA Issue: [IMPALA-10349](#)

Catalog and Event Processor Improvements

- Faster Inserts for Partitioned Tables (IMPALA-14051): Inserting data into very large partitioned tables is now much faster. Previously, Impala communicated with the Hive Metastore (HMS) one partition at a time, which was a major slowdown. Impala now uses the batch insert API to send all insert information to the HMS in one highly efficient call, significantly boosting the performance of your `INSERT` statements into transactional tables.
- Quicker Table Administration (IMPALA-13599): Administrative tasks, such as running `DROP STATS` or changing the `CACHED` status of a table, are now much faster on tables with many partitions. Impala previously made thousands of individual calls to the HMS for these operations. The system now batches these updates, making far fewer calls to the HMS and speeding up these essential administrative commands.
- Reliable Table Renames (IMPALA-13989): The `ALTER TABLE RENAME` command no longer fails when an `INVALIDATE METADATA` command runs at the same time. Previously, this caused the rename to succeed in the Hive Metastore but fail in Impala's Catalog Server. Impala now includes automatic error handling that instantly runs an internal metadata refresh if the rename is interrupted, ensuring the rename completes successfully without requiring any manual user steps.
- Efficient Partition Refreshes (IMPALA-13453): Running `REFRESH <table> PARTITION <partition>` is now much more efficient. Previously, this command always fully reloaded the partition's metadata and column statistics, even if the partition was unchanged. Impala now checks if the partition data has changed before reloading, avoiding the unnecessary drop-add sequence and significantly improving the efficiency of partition metadata updates.
- Reduced Partition API Calls (IMPALA-13599): Impala has reduced unnecessary API interactions with the HMS during table-level operations. Commands like `ALTER TABLE . . . SET CACHED/UNCACHED` or `DROP STATS` on large tables previously generated thousands of single `alter_partition()` calls. Impala now utilizes the HMS's bulk-update functionality, batching these partition updates to drastically reduce the total number of required API calls.
- `REFRESH` on multiple partitions (IMPALA-14089): Impala now supports using the `REFRESH` statement on multiple partitions within a single command, which significantly speeds up metadata updates by processing partitions in parallel, reduces lock contention in the Catalog service, and avoids unnecessary increases to the table version. See [Impala refresh](#)

- Impala cluster responsiveness during table renames(IMPALA-13631): This ensures that the critical internal lock is no longer held during long-running external calls initiated by ALTER TABLE RENAME operations. This prevents the entire Impala cluster from being blocked, allowing other queries and catalog operations to proceed without interruption.

Apache Jira: [IMPALA-14051](#), [IMPALA-13599](#), [IMPALA-13989](#), [IMPALA-13453](#), [IMPALA-14089](#) , [IMPALA-13631](#)

Enable global admission controller

It is now a standalone service that maintains a consistent view of cluster resource usage and can make admission decisions without risking over-admission. For more information, see [Impala components](#)

Auto-optimized parquet collection queries

Impala now automatically boosts query performance for tables with collection data types by setting the `parquet_late_materialization_threshold` to 1 when data can be skipped during filtering. This ensures maximum efficiency by reading only the data needed.

For more information, see [Impala lazy materialization](#)

Apache Jira: [IMPALA-3841](#)

Impala-shell now shows row count and elapsed time for most statements in HiveServer2 mode

When running Impala queries, some commands over HiveServer2 protocol (like REFRESH or INVALIDATE) did not show the Fetched X row(s) in Ys output in Impala-shell, even though Beeswax protocol showed them.

This issue was resolved by adding a new option in Impala-shell called `--beeswax_compat_num_rows`. When this option is enabled, Impala-shell now prints Fetched 0 row(s) in along with the elapsed time for all Impala commands. This requires impala-shell 4.5.1a1 or higher. See, [Impala shell options](#)

Apache Jira: [IMPALA-13584](#)

Support for arbitrary encodings in text and sequence files

Impala now supports reading from and writing to Text and Sequence files that use arbitrary character encodings, such as GBK, beyond the default UTF-8.

Impala now recognizes the Hive table property "serialization.encoding" used with LazySimpleSerDe.

- When reading data from these files, Impala uses the specified encoding to correctly decode the string data (example, converting GBK bytes into readable characters).
- When inserting data into these tables, Impala correctly encodes the inserted strings into the specified charset before saving them to the text file.

For more information, see [Impala textfile](#) .

Apache Jira: [IMPALA-10319](#)

New query options for reliable metadata synchronization

Impala now offers new query options to give you a reliable way to ensure your queries run with the latest table data after the relative HMS modifications are done.

- `SYNC_HMS_EVENTS_WAIT_TIME_S`: Sets the maximum time, in seconds, you are willing to wait for the metadata synchronization from the Hive Metastore (HMS). Setting this enables Impala to pause query compilation automatically until changes are applied, ensuring metadata consistency.
- `SYNC_HMS_EVENTS_STRICT_MODE`: Controls error handling if the wait time is exceeded. By default, Impala proceeds with a warning. Set to TRUE to force the query to fail immediately, guaranteeing strict consistency.

See, [Impala query options](#)

JIRA Issue: [IMPALA-12152](#)

Impala now supports Hive's legacy timestamp conversion to ensure consistent interpretation of historical timestamps

When reading Parquet or Avro files written by Hive using legacy timestamp conversion, Impala's timezone calculation for UTC timestamps could be incorrect, particularly for historical dates and timezones like Asia/Kuala_Lumpur or Singapore before 1982. This meant the timestamps displayed in Impala were different from those in Hive.

This issue is addressed by Impala now checks for the `writer.zone.conversion.legacy` flag in the Parquet file metadata to determine if Hive's legacy timestamp conversion method was used. If found, Impala uses a compatible conversion method. For older files without this flag, a new session and query option, `use_legacy_hive_timestamp_conversion`, has been added to control the conversion method. See, [Impala query options](#)

Apache Jira: [IMPALA-13627](#)

Supporting one-dimensional arrays in Kudu tables

Impala now supports one-dimensional arrays in Kudu tables. You can create Kudu tables with array columns and perform selection queries on these complex collection types.

Impala now supports tables created with the Hive JDBC Storage handler

Impala now translates and adapts the hive JDBC table properties upon loading, making them compatible with Impala's internal requirements. Previously, Impala had difficulty reading tables created using the hive JDBC Storage handler due to differences in how table properties, such as JDBC driver and DBCP configurations, were defined compared to Impala-created tables. See, [Impala External JDBC Tables \(Preview\)](#)

Apache Jira: [IMPALA-12992](#)

New catalogd flag to disable HMS sync by default

You can now use the `disable_hms_sync_by_default` `catalogd` startup flag to set a global default for the `impala.disableHmsSync` property. This feature allows you to skip event processing for all databases and tables by default while opting in specific elements as needed.

For more information, see: [disable_hms_sync_by_default Options for catalogd Daemon](#)

Apache Jira: [IMPALA-14131](#)

Parallel metadata loading in local catalog mode

Previously, when a query accessed multiple unloaded tables in local catalog mode, Impala loaded the metadata for those tables one after another. This sequential process caused significant latency and performance regressions compared to the legacy catalog mode.

This issue is addressed by parallelizing the table loading process. The fix allows Impala to load and gather metadata for multiple tables simultaneously. You can control the maximum number of threads used for this process by using the new `max_stmt_metadata_loader_threads` flag, which defaults to 8 threads per query compilation. See,

[Catalog startup flag](#) Apache Jira: [IMPALA-14447](#)

Specifying compression levels for LZ4, ZLIB, and ZSTD

You can now specify compression levels for the LZ4, ZLIB, GZIP, and ZSTD codecs to achieve higher compression ratios. This includes support for high compression modes in LZ4 (levels 3–12) and negative compression levels for ZSTD. These levels are supported by using the `compression_codec` query option.

For more information, see [compression_codec query option](#)

Apache Jira: [IMPALA-10630](#), [IMPALA-14082](#)

Impala now supports ARM architecture

What's New in Apache Kafka

New features and functional updates for Kafka are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces new features of Kafka and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see [New Features](#).

Rebase on Kafka 3.9

Kafka shipped with this version of Cloudera Runtime is based on Apache Kafka 3.9.1 (previously 3.4.1). For more information, see the following resources:

- Notable changes for releases 3.5.0 through 3.9.1: [Upgrading | Apache Kafka](#).
- The Apache Kafka release notes for the following versions:
 - [Kafka 3.5.0](#)
 - [Kafka 3.5.1](#)
 - [Kafka 3.6.0](#)
 - [Kafka 3.6.1](#)
 - [Kafka 3.5.2](#)
 - [Kafka 3.7.0](#)
 - [Kafka 3.6.2](#)
 - [Kafka 3.7.1](#)
 - [Kafka 3.8.0](#)
 - [Kafka 3.8.1](#)
 - [Kafka 3.9.0](#)
 - [Kafka 3.9.1](#)
- The Apache Kafka release announcements: [Release Announcements | Apache Kafka](#)

KRaft is generally available and ZooKeeper is deprecated

KRaft (Kafka Raft) is generally available. KRaft is from now on the recommended metadata management mode for Kafka in Cloudera. Additionally, migrating existing ZooKeeper-based Kafka clusters to use KRaft is now possible.

With the general availability of KRaft, deploying new or using existing Kafka clusters running in ZooKeeper mode is deprecated. Additionally, support for ZooKeeper-based Kafka clusters will be removed in a future release.

Cloudera recommends the following:

- Deploy all new Kafka clusters in KRaft mode.
- Migrate existing ZooKeeper-based clusters to KRaft following an upgrade to Cloudera Runtime 7.3.2.

This is the only version where migration is possible. Neither previous or future major, minor, and maintenance versions support migration.

For additional information, see the following resources:

- [Installing Kafka in Cloudera Base on premises](#)
- [Kafka KRaft](#)
- [Migrating Kafka from ZooKeeper to KRaft overview](#)

Kafka protocol and metadata version is set automatically during upgrades

When upgrading Kafka, Cloudera Manager now automatically sets the `inter.broker.protocol.version` property for ZooKeeper-based clusters and the `metadata.version` property for KRaft-based clusters. You no longer need to manually set these properties to the current protocol or metadata version before an upgrade. This feature is only available when upgrading to Cloudera Runtime 7.3.2 or higher.

After the upgrade, clearing these properties remains a manual task. However, in Cloudera Runtime 7.3.2 and higher, both `inter.broker.protocol.version` and `metadata.version` are now available for direct configuration in Cloudera Manager Kafka Configuration . The label names of the properties are Kafka Inter-Broker Protocol Version and Kafka Metadata Version. This means you can set or clear these properties directly from the UI, without needing to use advanced configuration snippets.

Connector-level offset flush control

A new connector-level property, `cloudera.offset.flush.interval.ms`, is added. Use this property to override the Kafka Connect role-level Offset Flush Interval (`offset.flush.interval.ms`) property. Overriding enables you to control the interval at which connector task offsets are committed on a per-connector basis.

Configure `cloudera.offset.flush.interval.ms` in connectors that need a different offset flush interval than the role default. This is commonly useful for connectors where the interval controls how often data is flushed to target systems, for example `NiFiStatelessSink`, `HDFSSink`, and `S3Sink`.

IPv6 support for Kafka

Starting with the 7.3.2 release, Kafka supports IPv6 with dual-stack functionality, allowing seamless communication over both IPv4 and IPv6 networks. This capability improves network scalability, future-proofs deployments, and enhances overall platform security.

Offline Log Directories chart

A new default chart, **Offline Log Directories**, is added for Kafka in Cloudera Manager. This chart can help you quickly identify and track storage issues on your brokers. It is available by default for the Kafka service as well as for individual Kafka Broker role instances.

The chart shows offline log directories and their mount paths for Kafka brokers. A non-zero value indicates an active error state for a specific log directory, while a value of 0 means the directory was in an error state during the selected timeframe but is now healthy. The chart only displays log directories that had errors during the selected timeframe.

New actions for collecting Kafka diagnostic data

The following new service-specific actions are available for collecting Kafka diagnostic data in Cloudera Manager:

- Collect Kafka Cluster Diagnostics - gathers detailed cluster-wide data, including topics, configurations, consumer groups, and more.
- Describe Kafka Topics - provides detailed information about all Kafka topics.

These actions are available in the Actions dropdown on the Kafka service and Kafka Broker role instance pages. Diagnostic data is printed to stdout for immediate access and also saved as a compressed archive on the host where the action runs.

For more information, see [Collecting Kafka diagnostic data using Cloudera Manager actions Connect](#).

Debezium connectors upgraded from 1.9.8.Final to 3.3.1.Final

This release of Cloudera Runtime ships version 3.3.1.Final of the following Debezium connectors:

- MySQL
- PostgreSQL
- Oracle
- SQL Server

- Db2

Existing connector instances are automatically upgraded to the new version as part of a cluster upgrade. However, you will be required to make configuration updates before you can upgrade your cluster. Critical changes that affect all Debezium connectors are summarized below.



Important: Debezium 3.3.1.Final includes many major breaking changes compared to 1.9.8.Final. Depending on your connector configuration and use case, additional configuration updates not highlighted here might be necessary. For a comprehensive list of changes, review release notes on <https://debezium.io/>.

- Property renaming (configuration namespace changes)

New, more consistent namespaces for configuration properties are introduced. The old `database.*` prefixes have been removed. Connector configuration keys collected in the following table must be updated before an upgrade.

| Old Property Prefix (Debezium 1.9) | New Property Prefix (Debezium 3.3) |
|---|--|
| <code>database.server.name</code> | <code>topic.prefix</code> |
| <code>database.history.*</code> | <code>schema.history.internal.*</code> |
| <code>database.*</code> (JDBC pass-through) | <code>driver.*</code> |
| <code>database.dbname</code> (SQL Server) | <code>database.names</code> |

- Database driver version requirements are updated

The recommended and supported JDBC driver versions used by the majority of connectors has changed. The following table collects the JDBC drivers you will need to deploy on your cluster before an upgrade.



Note: The recommended driver version for Debezium Db2 remains unchanged and is still 11.5.0.0 but is listed in the following table as a reference.

| Component | New Driver Version / Notes |
|------------|---|
| MySQL | 9.1.0 |
| PostgreSQL | 42.7.7 |
| Oracle | 21.x, 23.x — use a Java 11+ Oracle JDBC driver (<code>ojdbc11.jar</code>) |
| SQL Server | 12.4.2.jre8 |
| Db2 | 11.5.0.0 |

For step-by-step upgrade instructions on configuration updates and driver replacements, see [Getting Started Upgrading a Cluster](#).

What's New in Apache Knox

Learn about the new features of Apache Knox in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.0:

Cloudera Runtime 7.3.2 introduces new features of Knox and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see [New Features](#).

SameSite attribute for pac4j session cookies is now configurable

You can now configure the SameSite attribute for pac4j session cookies.

Group impersonation support in Knox

Knox now supports group impersonation, allowing users in specific groups to impersonate other users. For more information, see [Configuring Group Impersonation in Knox](#).

Knox IDBroker integration with HashiCorp Vault

Knox IDBroker now integrates with HashiCorp Vault for AWS credentials management, allowing IDBroker to authenticate with AWS using short-lived credentials from Vault instead of storing long-lived credentials for this purpose. For more information, see [Configuring Knox IDBroker with HashiCorp Vault](#).

Role-level alias management for Knox Gateway and IDBroker

The alias management configuration has been moved from service-level to role-level. Each role now has its own dedicated configuration: `gateway_save_alias_command_input` for the Knox Gateway role and `idbroker_save_alias_command_input` for the IDBroker role. Two role-specific commands are now available: Save Alias - Knox Gateway and Save Alias - IDBroker. For more information, see [Saving aliases](#).

What's New in Apache Kudu

New features and functional updates for Kudu are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces new features of Kudu and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see [New Features](#).

Kudu JDK 17 upgrade

Kudu supports JDK 17. Upgrading to JDK 17 ensures your application stays competitive and maintainable by providing improved performance, increased security, modern language features, and long-term support.

IPv6 support for Kudu

Starting with the 7.3.2 release, Kudu supports IPv6 with dual-stack functionality, allowing seamless communication over both IPv4 and IPv6 networks. This capability improves network scalability, future-proofs deployments, and enhances overall platform security.

For more information, see [Enable Dual-Stack support for Kudu](#).

Support for one-dimensional arrays in Kudu

Kudu now supports one-dimensional arrays for scalar types, such as integers, strings, booleans, and binaries. This feature allows you to store multiple related values in a single column, which improves query performance and simplifies your data model by removing the need for complex joins or denormalized tables.

For more information, see [Manage one-dimensional arrays in Kudu](#).

Replicating Kudu tables using Apache Flink

You can now use the Apache Flink-based replication job to continuously replicate data from a source Kudu cluster to a destination Kudu table. This feature uses Kudu diff scans to efficiently capture and move only the changed rows, supporting high-availability and disaster recovery scenarios.

For more information, see [Replicating Kudu tables using Apache Flink](#).

What's New in Livy

There are no new features in Livy in Cloudera Runtime 7.3.2.0.

What's New in Navigator Encrypt

Learn about the new features of Navigator Encrypt in Cloudera Runtime 7.3.2.0, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.0:

Navigator Encrypt installed using Cloudera Manager

In Cloudera Runtime 7.3.2.0 and higher releases, the new Navigator Encrypt parcel centralizes the entire lifecycle management within Cloudera Manager. Previously, Navigator Encrypt required manual, host-by-host RPM installation and configuration. Administrators can now manage the installation, configuration, and upgrade of Navigator Encrypt across all cluster hosts directly from the Cloudera Manager interface. This new approach provides a significant improvement in usability, streamlining ongoing maintenance and lifecycle management.

For more details, see *Navigator Encrypt installed using Cloudera Manager*.

What's New in Oozie

Learn about the new features of Oozie in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Oozie SSH action port is configurable

The Oozie SSH action port always uses the default port number 22. The default behaviour is still the same. But, you can change the port number using the following methods:

1. Use the new 0.4 schema version for your SSH action in your workflow.xml and add the port XML element with the new value.

Example

```
<port>11110</port>
```

2. Set the port number globally by adding the following oozie-site.xml safety-valve in Cloudera Manager with the new value.

Example,

```
"oozie.action.ssh.action.port: 11110"
```

If you set both the options, then the value set in the workflow.xml takes precedence and the port with the new port number should open in the target host.

What's New in Ozone

Learn about the new features of Ozone in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces new features of Ozone and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see [New Features](#).

Cloudera Storage Optimizer

Cloudera Storage Optimizer is an intelligent data lifecycle management feature that automatically reduces storage usage by converting infrequently accessed data from replicated storage (RATIS 3x)

to space-efficient erasure coding (EC). It analyzes access patterns and applies configurable policies to identify and convert cold data, reducing the storage footprint by 45% to 60% while maintaining data durability and availability. The benefit is additional usable capacity for the same licensed capacity. For more information, see [Cloudera Storage Optimizer](#) documentation.

Storage container reconciliation

Generally, container replicas are identical in nature and contain the same data. Container reconciliation is a repair mechanism that allows administrators to detect container replicas that are corrupted or have otherwise diverged from each other for any reason, and repair these replicas, making them identical again. The data within each replica is summarized as a single data checksum. When the data within the replicas diverges, the replicas will have different data checksums. After reconciling the replicas, they will each contain the same data and each replica's data checksums will match the others. This data checksum is designed to remain constant even as data is deleted from the container in the background. For more information, see [Storage container reconciliation](#) documentation.

Support for G1 Garbage Collector (G1GC) with JDK 17

When running Ozone on JDK 17, all Ozone server processes now default to the G1 Garbage Collector (G1GC). No action is required for most deployments.

Configuration Changes for the FIPS-compliant SASL Changes

Cloudera supports DIGEST-SHA, a new SASL mechanism, to replace DIGEST-MD5. DIGEST-SHA is similar to DIGEST-MD5 except that DIGEST-SHA uses SHA256 and AES instead of MD5 and DES for message digest and encryption, respectively. As a result, DIGEST-SHA uses only FIPS-compliant algorithms. For more information, see [Step 1: Prepare hosts](#).

Moved the port `ozone.prometheus.http-port` from 9094 to 9096

The port `ozone.prometheus.http-port` is moved from 9094 to 9096 to avoid Ozone Prometheus port colliding with HBase thrift server if they are on the same host.

Added new Ozone canaries

To configure the threshold value, the following new Ozone canaries are added along with the relevant configuration property that can be used to disable them:

- Ozone SCM Heap memory canary: `ozone_scm_heap_canary_enabled`
- Ozone OM Heap memory canary: `ozone_om_heap_canary_enabled`
- Ozone Datanode failed volumes canary: `ozone_datanode_failed_volume_canary_enabled`

Added role-specific Java options

Previously, Ozone Java options (JAVA_OPTS) were configured through a single shared setting. This change introduces separate Java option parameters for each Ozone service role, allowing more granular configuration.

The following role-specific parameters are now available:

- Storage Container Manager (SCM): `ozone_scm_java_opts`
- Ozone Manager (OM): `ozone_om_java_opts`
- DataNode (DN): `ozone_dn_java_opts`
- S3 Gateway (S3G): `ozone_s3g_java_opts`
- Recon: `ozone_recon_java_opts`

Each parameter applies only to its corresponding service and overrides the shared JAVA_OPTS configuration where applicable.

Added Knox proxy user configuration for Ozone Recon

Following new configuration properties to control Knox proxy user (impersonation) access to Ozone Recon are introduced:

- `ozone.recon.http.auth.proxyuser.knox.hosts`: Specifies the hosts from which Knox is allowed to impersonate users when accessing Recon.
- `ozone.recon.http.auth.proxyuser.knox.groups`: Specifies the user groups that Knox is allowed to impersonate when accessing Recon.

These settings restrict which hosts and user groups are permitted for Knox-based impersonation in Recon.

Added new verifier for container states

The `ozone debug replicas verify` command now includes a new verifier: `--container-state`.

```
ozone debug replicas verify --container-state <VOLUME/BUCKET/PATH> -o <OUTPUT.JSON>
```

This new verifier checks the container state in the Storage Container Manager (SCM) for every container in a replica. It identifies keys mapped to containers in problematic states, such as:

- Containers: DELETED or DELETING
- Replicas: DELETED, UNHEALTHY, or INVALID

The command generates a JSON file containing details on keys, blocks, and replicas, along with a pass/fail status for each check. It supports the following options:

- `--all-results`: Displays all verification outcomes
- `-o`: Specifies the output file path
- `--container-cache-size`: A new option that sets the number of containers stored in the in-memory cache for the `--container-state` check.

The default value is 1 million containers (approximately 43MB of memory). The value must be greater than zero; otherwise, the system uses the 1 million default.



Important: The system ignores this option if you do not use the `--container-state` flag.

Enhanced volume health checks triggered by container scanner

When a container scanner (background or on-demand, data or metadata) marks a container as unhealthy due to corruption, the system now automatically triggers an on-demand scan of the associated volume. This enhancement helps detect and address broader volume-level issues that may exist beyond the affected container, improving overall data integrity and reliability.

New container health metrics

Starting from Cloudera Base on premises 7.3.2.0 release, Ozone Recon now exposes additional metrics for the Container Health Task to improve observability of container state across the cluster. Recon tracks the counts of:

- Missing containers
- Under-replicated containers
- Over-replicated containers
- Mis-replicated containers

Increased default Ratis segment size for Ozone Manager (OM) and Storage Container Manager (SCM)

To improve startup performance, the default values of `ozone.om.ratis.segment.size` and `ozone.scm.ha.ratis.segment.size` configurations are increased from 4 MB to 64 MB.

- The `ozone.om.ratis.segment.size` configuration sets the size of the raft segment used by Apache Ratis on OM.
- The `ozone.scm.ha.ratis.segment.size` configuration sets the size of the raft segment used by Apache Ratis on SCM.

Renaming all legacy HDFS configuration keys and the corresponding Java constants to HDDS

All legacy configuration keys previously prefixed with `dfs.` (used by HDFS) have been updated to use the `hdds.` prefix, aligning them with the HDDS component. Backward compatibility is maintained. The system will continue to recognize the old `dfs.` keys, but they are now deprecated. Deprecation notices and mappings are handled internally by adding `DeprecationDelta` in `OzoneConfiguration`.

Additionally, the corresponding Java constants for both the configuration key and the default value are renamed accordingly.

All deletion configurations are dynamically configurable without restart

Starting from Cloudera Base on premises 7.3.2.0 release, the following deletion configurations are dynamically reconfigurable without requiring a restart:

- For Ozone Manager: `ozone.thread.number.dir.deletion`
- For DataNode:
 - `ozone.block.deleting.service.interval`
 - `ozone.block.deleting.service.timeout`

Support callback on completed reconfiguration

The reconfiguration framework now supports a reconfiguration completed event. Components can now register callbacks that trigger after all properties in a task are processed. This allows the system to validate and apply interdependent settings automatically. The status and values of the reconfigured properties are accessible through `ReconfigurableBase#getReconfigurationTaskStatus` within the callback.

Moved and enhanced Ratis Log Parsing command in Ozone

The `ozone debug ratislogparser` command has been moved under the new `ozone debug ratis parse` subcommand. The `parse` subcommand now supports parsing different Ratis files, providing a more general and extensible interface.

A new optional flag `--role=[om, scm, datanode]` has been added to specify which schema should be used when parsing logs. By default, the schema is set to generic.

Enhanced output for `ozone debug replicas chunk-info` command

The `ozone debug replicas chunk-info` command now provides detailed information about the chunks associated with a given key. Previously, chunk-related details were only available when the `--verbose` flag was used. With this update, the command includes the `BlockData` retrieved from the `getBlock` call, along with additional details about each replica. The JSON output structure has also been improved for clarity.

Unified JSON output for `ozone debug replicas verify` command

The `ozone debug replicas verify` command now outputs JSON information for each key and the checks performed on it. As new verification checks are added, their results will be included in the same JSON objects. By default, the output is skipped for keys that pass all specified checks, reducing unnecessary output.

Increased Ozone Manager (OM) default snapshot limit

The default value for the Ozone Manager property `ozone.om.fs.snapshot.max.limit` has been increased, allowing support for up to 10,000 snapshots by default (previously 1,000). This change enables users to create and manage a significantly larger number of snapshots without requiring manual configuration changes.

Ozone Recon UI: Added new Cluster Health panel to the Grafana Dashboard

A new Cluster Health panel has been added to the existing Ozone - Overall Metrics Grafana dashboard. This panel displays the cluster growth rate, calculated dynamically using Prometheus metrics. By default, the panel shows the growth rate over the last 1 hour, but users can configure the time window as needed.

Renamed Ozone configuration property `ozone.compaction.service.enabled`

The configuration property `ozone.compaction.service.enabled` has been renamed to `ozone.om.compaction.service.enabled`. This change aligns with Ozone's configuration key conventions, as the property is used only by the Ozone Manager (OM).

Added bucket layout flag to `freon rk` command

Previously, the `freon randomkeys (rk)` command did not allow users to specify the bucket layout, always creating buckets with the default layout. With this improvement, a new option has been introduced to set the bucket layout through the command, providing greater flexibility.

Enhanced DataNode listing with usage-based sorting

The `ozone admin datanode list` command now supports two new CLI options: `--most-used` and `--least-used`. These options enable users to sort datanodes based on their usage. When used, the command output includes additional details such as used space, total capacity, and percentage of space used for each datanode. This enhancement makes it easier to identify the most and least utilized datanodes in your Ozone cluster.

Updated the Ozone SCM Container Placement policy description

The descriptions for the following Ozone configuration properties have been corrected in the `ozone-default.xml` to clarify that the default placement policy is `SCMContainerPlacementRackAware`, not `SCMContainerPlacementRandom`:

- `ozone.scm.container.placement.impl`
- `ozone.scm.container.placement.ec.impl`

The code fallback for an unset `ozone.scm.container.placement.impl` has also been updated to use `SCMContainerPlacementRackAware`, similar to the `ozone-default.xml`.

Enhanced container creation command with replication options

The `ozone admin container create` command has been improved to support additional replication configurations. Previously, containers could only be created with a replication factor of ONE and a replication type of STANDALONE. Administrators can now specify replication types (RATIS and EC), replication factors, and algorithms when creating containers.

```
ozone admin container create --replication=<REPLICATION> --replication-type=<TYPE>
```

What's New in Apache Phoenix

New features and functional updates for Phoenix are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2:

IPv6 support for Phoenix

Starting with the 7.3.2 release, Phoenix supports IPv6 with dual-stack functionality, allowing seamless communication over both IPv4 and IPv6 networks. This capability improves network scalability, future-proofs deployments, and enhances overall platform security.

For more information, see [Enabling IPv6 support for Phoenix](#).

Phoenix supports cluster-wide metadata upgrade blocks

Phoenix now provides the ability to block cluster-wide metadata auto-upgrades. To prevent an upgrade from occurring, you must add a `BLOCK_UPGRADE` row in the `SYSTEM.MUTEX` or `SYSTEM:MUTEX` HBase table (depending on whether namespace mapping is enabled), in the `0:MUTEX_VALUE` column family, before you start Phoenix with the upgraded versions. To add the row by using the HBase shell, run the following command:

```
put 'SYSTEM:MUTEX', 'BLOCK_UPGRADE', '0:MUTEX_VALUE', 'MUTEX_LOCKED'
```

For a Java code example, see the `LoadSystemTableSnapshotBase.java` file in the [Apache Phoenix repository](#).

Phoenix JDK 17 upgrade

Phoenix supports JDK 17 starting with the 7.3.2 release; upgrading to JDK 17 ensures your application stays competitive and maintainable by providing improved performance, increased security, modern language features, and long-term support.

Deprecation of OMID service

The OMID service is removed from the Cloudera on cloud installation of Cloudera Runtime 7.3.2. This change is implemented because OMID appears to be unused in these environments. Consequently, the OMID service is no longer automatically deployed, effective immediately.

Users who still require OMID are not losing access to the service. You can still install OMID manually through Cloudera Manager if needed for your operations.

Apache Phoenix component upgraded to 5.2.1

The Phoenix runtime component upgrades from 5.1.3 to 5.2.1. This release delivers performance enhancements and introduces cluster-wide metadata upgrade blocks, providing you with greater control during upgrades and increasing stability in modern environments.

What's New in Apache Ranger

Learn about the new features of Apache Ranger in Cloudera Runtime 7.3.2.0, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.0:

The new features for Apache Ranger in Cloudera Runtime 7.3.2.0 include cumulative updates from previous releases. This version specifically incorporates features introduced in Cloudera Runtime 7.3.1.100 through 7.3.1.706. For a complete list, see [New Features](#).

Support for Amazon S3-compatible object stores through RAZ

Support has been added for Amazon S3-compatible object stores through the Ranger Remote Authorization Service (RAZ).

Cloudera Base on premises defaults to HDFS and natively supports the Ozone object store. There was limited support for Amazon S3-compatible object stores with respect to executing workloads like Hive and Spark. Not all workloads were supported from an authorization perspective on Amazon S3-compatible object stores. The inconsistent policy framework created challenges for the Policy Administrator in terms of applying corporate policies to secure sensitive data. Managing data access across teams and individuals was an architectural challenge.

The Ranger RAZ resolves this challenge by using Apache Ranger policies to authorize access to Amazon S3-compatible object storage, similar to HDFS files. For details, see [Access Control for Amazon S3-compatible object stores](#).

Ranger has been upgraded

Ranger has been upgraded from version 2.4.0 to 2.6.0, incorporating the latest fixes and improvements from the Ranger project while maintaining compatibility with existing Cloudera deployments.

Ranger HBase plugin optimizations

You can now enable column authorization optimization, wherein column authorization is skipped when a user is fully authorized for the column families. To enable this optimization, add the following Ranger service configuration for `cm_hbase` from Ranger UI: `ranger.plugin.hbase.column.auth.optimized=true`. The new configuration can lead to significantly better performance of multiget and multiput workloads, wherein thousands of columns may be accessed together. No HBase service restart is required to enable or disable this configuration.

When the optimization is enabled, there is a behavior change in auditing in Ranger. For more information, see [Behavioral Changes in Apache Ranger](#).

Ranger RMS supports for multiple storage types

Starting with Cloudera Base on premises 7.3.2.0 release, RMS supports S3 storage in addition to HDFS and Ozone, allowing it to track mapping data for all these storage types simultaneously.

After the first full-synchronization run, RMS downloads mappings for tables and databases present in the Hive Metastore (HMS). These tables/databases could be located in S3, OZONE, and HDFS file systems. RMS will map the tables and databases with their respective storage locations and store them in the Ranger database with their associated service ID (cm_s3, cm_ozone, or cm_hdfs). When raz-s3-chained-plugin, ranger-ozone-plugin, or ranger-hdfs-plugin requests mappings, it only retrieves the mappings relevant to its service. Specifically, raz-s3-chained-plugin downloads mappings exclusively for tables/databases whose storage location is S3, ranger-ozone-plugin for Ozone, and ranger-hdfs-plugin for HDFS.



Important: For users who had RMS configured for HDFS/Ozone before Cloudera Base on premises 7.3.2.0 and also have external table data stored in S3, a full synchronization is necessary after upgrading. This ensures that any databases and tables created prior to the upgrade, and whose data resides in S3, are correctly mapped. Any newly created tables or database mappings will be synchronized in RMS.

User's last login timestamp visibility

The Ranger Admin Web UI now displays the logged-in user's last login date and time on all pages.

Earlier, this information was only available under `Ranger Audits Login Sessions`.

Now the user's last login timestamp is persistently visible across all pages. For a user's initial login, the last login timestamp will not be displayed. Upon subsequent logins, the timestamp of the preceding successful login will be displayed.

Storage change for Ranger admin audit logs

Starting from Cloudera Runtime 7.3.2.0, Ranger admin audit logs are stored in the `x_trx_log_v2` table in the Ranger database. In releases earlier than 7.3.2.0, Ranger admin audit logs were stored in the `x_trx_log` table in the Ranger database. Hence, if you upgrade to Cloudera Runtime 7.3.2.0 from any previous version and you want to retain your old data, you must migrate the Ranger admin audit logs from the `x_trx_log` table to the `x_trx_log_v2` table.



Note: In the case of a fresh installation of Cloudera Runtime 7.3.2.0, this migration process is not required.

For more information on how to migrate the data, see [Migrating Ranger admin audit logs](#).

Back up ranger_audits Solr collection

When upgrading to Cloudera Runtime 7.3.2.0 and later releases that update the Ranger Audits Solr schema, Cloudera Manager runs an internal service command to apply the schema changes safely.

Although the command is designed to be non-destructive, Cloudera strongly recommends creating a backup of the Ranger Audits Solr collection before starting the upgrade. This ensures a quick restoration of audit data if anything goes wrong during the upgrade.

For more information, see [Back up ranger_audits Solr collection](#).

Ranger mixed case group comparison

When Ranger Usersync is configured with case conversion and special character replacement using Regular Expression (regex), Ranger Usersync transforms the original user or group names from the source, for example, AD or LDAP, before storing them in the Ranger Admin database. Previously, if a Ranger plugin used the original name during authorization checks, the check failed because the Ranger Admin only recognized the transformed name.

This issue is now fixed. The fix is configurable at the plugin level using the `ranger.plugin.<serviceType>.supports.name.transformation` property, allowing users to enable or disable transformation based on their environment needs. For more information, see [Handling inconsistent username and group name conventions for consistent authorization](#).

ranger.plugin.hive.urlauth.filesystem.schemes configuration with default value of hdfs:

The default value of the `ranger.plugin.hive.urlauth.filesystem.schemes` configuration is `hdfs:,file:,wasb:,adl:`. When the `ranger.plugin.hive.urlauth.filesystem.schemes` configuration includes `hdfs:` as a default filesystem scheme in both HMS and HiveServer2, and you perform a CREATE EXTERNAL TABLE operation with a LOCATION clause, if both the Ranger Hive and HDFS plugins are enabled, the Ranger Plugin Authorizer performs authorization checks on all files within the target directory. This results in performance degradation that scales up linearly with the number of files.

The default value of the parameter remains unchanged, but a warning message has been added for the user. For more information, see [Hive authorizer URL policy with hdfs default value](#).

Added validation for whitespace in Ranger policy names

Ranger Admin now validates and handles leading and trailing whitespace in policy names and related properties (for example, `db_name` vs `db_name`) created through the Ranger API. Policies that contain unintended whitespace are detected and normalized so they no longer appear inconsistent between the API and the Ranger Admin web UI, helping prevent confusing behavior and reducing issues for application teams.

Bulk delete policies using policy name prefix

Added a Ranger REST API endpoint that supports bulk deletion of policies by policy name prefix (wildcard). Administrators can now delete multiple policies in a single request by specifying a service name and a policy name prefix (for example, `JAMES*`), improving operational efficiency over deleting policies one by one.

Support SASL bind for Ranger Usersync with AD/LDAP

Introduced support for SASL GSSAPI authentication in Ranger Usersync when connecting to AD/LDAP. This allows customers to avoid using simple bind credentials and mitigates performance and reliability issues observed with SSSD on RHEL 8.8 (for example, incomplete getent results with `enumerate=true`).

Ranger RAZ service-admin delegation support

Ranger RAZ now supports delegation for users with service-admin privileges. When making authorization calls to RAZ, users defined in the Ranger service configuration as `service.admin.users` or belonging to groups in `service.admin.groupscan` authorize access on behalf of other users. Requests from these service-admin users are also trusted for details such as the resourceOwner, improving support for administrative and service-level access workflows.

RAZ to support authorization for CAII service on Cloudera Base on premises

In Cloudera Base on premises 7.3.2.0, Ranger RAZ now supports the Cloudera AI Inference (CAII) service for authorization. During initialization, Ranger automatically creates the CAII service definition and a corresponding service instance, similar to other `RAZ#managed` services (for example, `cm_hive` and `cm_s3`). RAZ has been enhanced to authorize access to CAII resources and to honor `service.admin.users` and `service.admin.groups` so that designated service administrators can authorize access on behalf of other users.

Improved handling of Kafka super users in Ranger Kafka plugin

The Ranger Kafka authorizer now automatically parses the `Kafka.super.users` configuration property and adds those users to the plugin's super-user list. This allows Strimzi and other deployments that rely on super.users (for example, cluster operator and broker users) to have the expected unrestricted access to Kafka resources without requiring duplicate super-user configuration in Ranger.

Ranger Usersync LDAP referral handling improved

In Cloudera Runtime 7.3.2.0, the default value of the Ranger Usersync LDAP referral configuration has been changed from ignore to follow (`ranger.usersync.ldap.referral=follow`). This enables Ranger Usersync to follow LDAP referrals during group searches, which helps prevent `javax.naming.PartialResultException: Unprocessed Continuation Reference(s)` errors commonly seen with Active Directory LDAP and reduces the need for relying on the AD Global Catalog for referral resolution.

IP-based access control for policies

Ranger now supports IP-based access control across additional services by using context enrichers and policy conditions. Administrators can configure IP location data files, register the `RangerFileBasedGeolocationProvider` context enricher, and define IP-based policy conditions (for example, `accessed-from-ip`) to allow or deny access based on client IP ranges. This capability is not enabled by default and must be explicitly configured, following the steps described in the Apache Ranger Geolocation based policies documentation.

Improved REST API responses for `policy?policyNamePartial`

Ranger policy REST API responses for `policy?policyNamePartial` have been enhanced to optionally include policy metadata fields such as `createTime` and `updateTime`. These meta attributes are now configurable for retrieval, allowing administrators to access policy creation and modification timestamps directly from the public v2 policy API when enabled.

Ranger Swagger API usability improvement

The Ranger `/service/xusers/users` Swagger API now supports the optional query parameters `syncSource` and `userRole`. These parameters are exposed as fillable fields in Swagger to make it easier to construct complex user search requests (for example, filtering users by sync source such as Unix or LDAP/AD, or by role such as admin), especially for less technical users and environments with large user populations.

Improved `authzmigrator` script

Improved the `authz-export.sh` (`authzmigrator`) script so it no longer relies on a hard-coded Cloudera parcel path and JDBC driver. The script now uses the configured Cloudera Manager parcel directory (via `SENTRY_HOME`) and automatically selects the appropriate JDBC driver for the underlying database, preventing `NoClassDefFoundError: javax/jdo/JDOHelper` and making it more robust in environments with custom parcel layouts and non-default DB flavors.

Downlist the public group from Select Group option in New Policy page

To reduce the risk of accidentally assigning the Public group when creating a new Ranger policy, the Tab key no longer selects a group in the Select Group field on the New Policy page. Pressing Tab now only moves focus to the next field as expected, while the Enter key remains the standard way to confirm a group selection. This change significantly lowers the chance of unintended Public group assignment without impacting normal group selection behavior.

Delete Ranger RAZ access logs based on max number of files

To prevent Ranger RAZ access logs from filling up disk space on busy environments, Ranger RAZ now supports log rotation based on a configurable maximum number of access log files. Previously, access logs could only be rotated based on age (with a default of 15 days), which could still allow logs to exhaust the volume on high-activity clusters. With this enhancement, administrators can limit the number of retained Ranger Raz access log files, helping ensure Cloudera Manager and other services remain available even under heavy load.

Ranger access audit option to exclude resource

Improved Ranger access audit filtering to support excluding specific resources by name. You can now use an Exclude Resource Name filter in the Ranger Access audit UI so that audit searches can omit events for specified resources, leveraging existing backend Solr support for these exclusions.

Included `authz_export.tar.gz` for Sentry to Ranger migration with Cloudera Manager package

For CDH#to#CDP upgrades that involve Sentry#to#Ranger migration, the `authz_export.tar.gz` package is now bundled with the Cloudera Manager installation, in the same location as `authz-in-gest.zip`. This eliminates the need to contact Cloudera support to obtain the export archive from internal sources and simplifies authorization policy migration.

Ranger Admin logs to show the local timezone rather than UTC

Ranger Admin now records log timestamps in the local system timezone of the host instead of UTC. This aligns Ranger Admin logs with other Ranger components (such as UserSync and TagSync) and removes the need for manual log4j configuration, making troubleshooting and correlation with other system logs easier.

Cloudera Manager now sanitizes cluster names

Cloudera Manager now sanitizes cluster names that contain spaces when creating and evaluating Ranger plugin repositories that use the UNIQUE_PER_SERVICE naming strategy. Spaces in the derived Ranger repository (service) name are automatically replaced with underscores (_), ensuring compatibility with the Ranger constraint introduced in RANGER-2808 that disallows spaces in service names.

Improved Ranger Admin Diagnostic collection command from Cloudera Manager scripts

In Cloudera Manager 7.13.2.0, the Ranger Admin diagnostic collection command has been enhanced. A new configuration option, `ranger.admin.diag.metrics.collection.type`, allows you to control which types of Ranger metrics are collected during the Cloudera Manager diagnostics command. This helps avoid long-running Ranger Admin metrics collection (for example, cases where collection could exceed 40 minutes and cause the Cloudera Manager diagnostic command to stop) by letting you limit or skip specific metrics types as needed.

Updated Ranger RMS, TAGSYNC, and USERSYNC CSDs

Updated Ranger RMS, TAGSYNC, and USERSYNC Cloudera Service Descriptors (CSDs) to support the new Ranger RMS high availability (HA) implementation used by the `ranger-commo n-ha` module. The RMS server and `tagsync/usersync` service port configs are now managed as `first#class` CSD properties rather than via safety valves, improving HA configuration consistency and upgrade handling.

Cloudera Manager updates Ranger plugin service repositories with the correct cluster configurations

Cloudera Manager now updates Ranger plugin service repositories with the correct cluster configurations for test connection and resource lookup. For Ranger plugins on HBase, Knox, Atlas, and Ozone, Cloudera Manager automatically calculates and populates the required URLs, principals, and related settings instead of creating repositories with default, nonfunctional values. This reduces configuration errors and addresses frequent “Test connection” and lookup failures seen in customer deployments.

Enforce truststore configuration when Ranger plugin is enabled with Ranger admin SSL

When the Ranger plugin is configured to use Ranger Admin over SSL, Cloudera Manager now enforces truststore configuration for the plugin services. If the truststore settings are missing, the Ranger Plugin Services page displays a validation warning, helping prevent plugin sync failures caused by an unconfigured truststore.

Improved control of Ranger audit spooling to prevent disk exhaustion

Added Cloudera Manager configuration options for the Ranger plugins to control audit file spooling, including the spool file rollover interval and the maximum number of archived spool files. These settings are now exposed centrally for HDFS and Solr Ranger audits, reducing the risk of `/var/log` filling up when the audit destination is unavailable and eliminating the need for per-service `safety#valve` configuration.

Ranger audits collection creation reliability improved

Updated the Ranger audit configuration so that the `ranger.audit.solr.time.interval` setting now uses a 60-second (60000 ms) delay between retries when creating the `ranger_audits` collection in Solr, improving reliability in environments where Solr starts later than Ranger.

What's New in Ranger KMS

Learn about the new features of Ranger KMS in Cloudera Runtime 7.3.2.0, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.0:

Ranger KMS in a federated deployment

Ranger KMS can now be deployed in a federated cluster for key management.

A data cluster is a cluster where application and data processing occur. The cluster stores and processes actual datasets but does not directly manage encryption keys. A security cluster (with the Ranger KMS service installed) is managed separately from the data cluster. It manages the key lifecycle operations (for example, generation, rotation, storage). This separation of tasks enhances security by isolating the administration of the data and security clusters.

For more details, see *Installing Ranger KMS in a federated deployment*.

What's New in Schema Registry

New features and functional updates for Schema Registry are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

There are no new features introduced in this release.

What's New in Apache Solr

New features and functional updates for Solr are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2:

Solr JDK 17 upgrade

Solr supports JDK 17 starting with the 7.3.2 release. Upgrade to JDK 17 to gain improved performance, increased security, modern language features, and long-term support, ensuring your application remains competitive and maintainable.

What's New in Spark

Learn about the new features of Spark in Cloudera Runtime 7.3.2.0, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.0

Rebase Spark3 to Apache Spark 3.5.4 in Cloudera Runtime.

Spark 3.5.4 is the default Spark version in Cloudera Runtime. Refer to the [Apache Spark documentation](#) for more information on changes.



Important: Refer to [Migrating Spark applications](#) for more information on migrating your existing Spark applications.

Spark 3 is the default in Cloudera Runtime

Spark 3 is the default Spark version in Cloudera Runtime. Spark 2 has been removed and no longer available in 7.3.2.0



Important:

Spark 3 contains a large number of changes from Spark 2.

Refer to [Upgrading Spark](#) for more information on upgrading Spark clusters to 7.3.2.0, and [Migrating Spark applications](#) for more information on migrating your existing Spark applications between versions 2 and 3.

What's New in Spark Atlas Connector

There are no new features in Spark Atlas Connector in Cloudera Runtime 7.3.2.0.

What's New in Sqoop

New features and functional updates for Sqoop are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

There are no new features introduced in this release.

What's new in Streams Messaging Manager

New features and functional updates for Streams Messaging Manager are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

There are no new features introduced in this release.

What's New in Streams Replication Manager

New features and functional updates for Streams Replication Manager are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

Reverse Checkpointing

Streams Replication Manager now supports reverse checkpointing. This feature enables the tracking and replication of consumer offsets from a target cluster back to a source cluster. By tracking offsets in the reverse direction, you ensure that the progress made by consumer groups on a backup cluster is preserved and translated back to the primary cluster during a failback scenario.

Reverse checkpointing minimizes message duplication upon failback by mapping the offsets from the replica topic back to the equivalent offsets in the source topic. To enable this feature, you must configure the following in Cloudera Manager:

- Set the `cloudera.reverse.checkpointing.enabled` property to `true`.
- Enable bidirectional replication in the Streams Replication Manager's Replication Configs property.

In addition to service configurations, you must use the `srm-control` tool to explicitly allowlist topics for reverse checkpointing using the `reverse-checkpointed-topics` command. Consumer group replication must also be enabled in both directions.



Important: Reverse checkpointing relies on cluster prefixes to identify replica topics and is therefore not compatible with the `IdentityReplicationPolicy`.

Single REST server for all replication flows

Streams Replication Manager now uses a single REST server with a single port to handle inter-worker communication for all replication flows. Previously, a dedicated REST server was started for each replication flow. The new implementation exposes only the endpoints required for inter-worker coordination and task configuration updates. These endpoints are

restricted to inter-worker communication and cannot be accessed externally. The legacy per-flow REST server implementation is deprecated in 7.3.2 and will be removed in a future release. Cloudera recommends that you migrate your Streams Replication Manager clusters to the new implementation.

Suppressing internal metrics topics

You can now configure the Streams Replication Manager Service to suppress the eager creation of srm-metrics topics for all possible replication flows. This prevents the creation of unused topics. To enable this behavior, set the `metrics.topic.creation.for.possible.flows.enabled` property to false.

Configurable timeout for Streams Application Kafka Connection Health Test

A new SRM Service Streams Application Connection Test Timeout (`streams.replication.manager.service.streams.application.connection.test.timeout`) Cloudera Manager configuration option is now available for the Streams Replication Manager Service. It sets the timeout, in seconds, for the Streams Application Kafka Connection Health Test, which periodically checks connectivity to the target Kafka cluster. The default is 1 second.

What's New in YARN and YARN Queue Manager

New features and functional updates for YARN and YARN Queue Manager are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

Hadoop rebase summary

In Cloudera Runtime 7.3.2, Apache Hadoop is rebased to version 3.4.1. The Apache Hadoop upgrade improves overall performance and includes all the new features, improvements, and bug fixes from versions 3.2, 3.3, and 3.4.

Table 4: New features added from Apache Hadoop 3.2 to 3.4 versions

| Apache Hadoop version | Apache Jira | Name | Description |
|-----------------------|---------------------------|--|---|
| 3.3 | YARN-8762 | Interactive shell support in Hadoop YARN 3.3.0 | <p>Hadoop YARN 3.3.0 now incorporates the Interactive Shell for container execution. This feature provides the following capabilities:</p> <ul style="list-style-type: none"> Enables you to run commands and conduct real-time debugging within both standard YARN and Docker containers. Provides access through a new Command Line Interface (CLI) supporting the sh shell or using a dedicated terminal link within the new YARN UI2. Ensures that only the container owner is authorized to access the shell using stringent authentication protocols. Allows administrators to deactivate this feature using a specific configuration flag. |

Table 5: Improvements added from Apache Hadoop 3.2 to 3.4 versions

| Apache Hadoop version | Apache Jira | Name | Description |
|-----------------------|--------------------------------|--|---|
| 3.2 | MAPREDUCE-7461 | Added ability to specify user environment variables individually | Environment variables can now be specified for MapReduce tasks using individual properties. The following formats are used for specifying the environment variables: <ul style="list-style-type: none"> mapreduce.map.env.VARNAME=value mapreduce.reduce.env.VARNAME=value yarn.app.mapreduce.am.env.VARNAME=value yarn.app.mapreduce.am.admin.user.env.VARNAME=value |
| 3.3 | YARN-8489 | YARN Service Status Based on Dominant Component | The YARN service status reporting is now based on the status of the dominant component, improving the existing termination policies, for example, ALWAYS, NEVER. The final status of the dominant component determines the final state of the job. |
| 3.4 | YARN-9279 | YARN Hamlet Package Removal | The deprecated org.apache.hadoop.yarn.webapp.hamlet package is now completely removed to improve maintainability. This is an incompatible change in Hadoop YARN 3.4.0+. Applications relying on this old package must be updated to use the org.apache.hadoop.yarn.webapp.hamlet2 package. This affects the YARN webapp component. |
| 3.4 | YARN-10820 | Enhanced Reliability for YARN node list Command | The thread-safety issue is fixed in GetClusterNodesRequestPBImpl, that previously caused intermittent failures, such as java.lang.ArrayIndexOutOfBoundsException, with the YARN node list command. This change affects the YARN client in Hadoop YARN 3.4.0, 3.3.2, and 3.2.4, thereby, eliminating random crashes when running the YARN node list command. |

Table 6: Issues fixed between Apache Hadoop versions 3.2 to 3.4

| Apache Hadoop version | Apache Jira | Name | Description |
|-----------------------|--------------------------------|--|---|
| 3.3 | MAPREDUCE-7461 | MapReduce task initialization Timeout issue | Previously, MapReduce jobs stopped responding if a task terminated before sending its first heartbeat, as the task never timed out and remained stuck indefinitely in a "STARTING" state. This issue is now resolved by introducing a dedicated timeout mechanism specifically designed to catch and terminate tasks that fail to initialize and send their first heartbeat. |
| 3.4 | YARN-9809 | Miscommunication between RM and NM when NodeManagers are unhealthy | Previously, if a NodeManager (NM) was registered in an unhealthy state, it did not communicate the status immediately. As a result, the Resource Manager (RM) mistakenly scheduled many containers to that unhealthy node before the first heartbeat was received. Once the first heartbeat finally arrived, the RM recognize the unhealthy status and abruptly ended all the recently scheduled containers, causing unnecessary task failures and wasted resources. This issue is now resolved and NMs now explicitly supply their health status during their initial registration with the RM. |

What's New in Apache ZooKeeper

New features and functional updates for ZooKeeper are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

IPv6 support for ZooKeeper

Starting with the 7.3.2 release, ZooKeeper supports IPv6 with dual-stack functionality, allowing seamless communication over both IPv4 and IPv6 networks. This capability improves network scalability, future-proofs deployments, and enhances overall platform security.

ZooKeeper JDK 17 upgrade

ZooKeeper supports JDK 17 starting with the 7.3.2 release. Upgrade to JDK 17 to gain improved performance, increased security, modern language features, and long-term support, ensuring your application remains competitive and maintainable.

IPv6 prefixes support for ZooKeeper ACLs

Support for IPv6 prefixes is now added to ZooKeeper ACLs, enabling more flexible access control in IPv6-enabled deployments.

ZooKeeper component upgraded to 3.8.5

ZooKeeper is now upgraded to 3.8.5. This release incorporates stability and correctness fixes from the upstream, which deliver performance and maintenance improvements, and enhanced client connectivity support.

Fixed Issues In Cloudera Runtime 7.3.2

Fixed issues represent issues reported by Cloudera customers that are addressed in this release

Fixed Issues in Atlas

Fixed issues for Atlas are addressed in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2 resolves Atlas issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.700. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

CDPD-82058: UI displays current time instead of "NA" when server returns date fields as '0'

7.3.2

When the Atlas server returns date field values as 0 (zero) in API responses, the user interface incorrectly displays these fields as the current system time instead of showing "NA" (Not Available) to indicate that no valid date is present. This causes confusion when viewing entity details, audit logs, or other date-related information where timestamps are expected but not available. The fix modifies the UI utility functions in both the classic and React-based user interfaces to properly detect zero date values and display "NA" instead of converting them to the current time.

Apache JIRA: [ATLAS-5015](#)

CDPD-84505: Advanced Search not working properly in Atlas UI

7.3.2

When using the Advanced Search feature in the Atlas user interface, selecting a type definition (typeDef) and clicking the search button fails to return any results or fetch entities for the selected type. This issue affects the search result layout view in both the classic and React-based user interfaces. The broken functionality forces users to rely on Basic Search instead of the more powerful Advanced Search capabilities. The fix corrects the search result handling logic to properly process and display search results for type-based queries.

CDPD-87670: Atlas Glossary becomes unresponsive when page size is set to 50

7.3.2

When viewing a glossary item that has more than 25 associated entities and changing the page limit to 50 in the Glossary section, the Atlas user interface becomes unresponsive and displays an error alert with the message "expected type AtlasGlossaryCategory; found AtlasGlossaryTerm".

This issue was caused by the UI calling an incorrect API endpoint. The fix corrects the API call to properly handle larger page sizes in the Glossary section.

Apache JIRA: [ATLAS-5067](#)

CDPD-82062: UI displays current time instead of "NA" when server returns date fields as '0'

7.3.2

The fix updates date handling utilities in both the classic and React-based Atlas user interfaces so that when API responses contain date field values as 0 (zero), the UI displays "NA" (Not Available) instead of converting them to the current system time.

Apache JIRA: [ATLAS-5015](#)

CDPD-89296: Basic search API fails to validate nested and grouped filter criteria

7.3.2

Previously, the `/v2/search/basic` API endpoint validated only top-level, non-grouped filter criteria. When entity or tag filters used grouped conditions (created with the Add Group Filter option in the Atlas UI), the API did not recursively validate nested `FilterCriteria` objects. This caused requests with composite or grouped filters to fail with a bad request error such as "Invalid operator specified for attribute: null". The fix adds recursive validation that correctly handles both grouped filter conditions (which require an AND/OR condition and contain nested criteria) and leaf-level filter criteria (which require a valid operator, a non-blank attribute name, and a non-blank attribute value, except when using `IS_NULL` or `NOT_NULL` operators). Requests that fail validation now return a descriptive error message.

Apache JIRA: [ATLAS-5053](#)

CDPD-81913: Auto-created Kafka topics are not audited to Atlas on ZooKeeper-based clusters

7.3.2

When a Kafka topic was automatically created (for example, following a produce attempt on a non-existent topic), Kafka set the error code to 3 (`UNKNOWN_TOPIC_OR_PARTITION`) in the `MetadataResponse`. The Atlas Kafka hook uses an entity selector that only processes audit events with a zero error code, causing topic auto-creation events to be silently dropped and not audited to Atlas. This issue only affected ZooKeeper-based Kafka clusters; KRaft-based clusters handle topic creation differently and were not affected. The fix updates the Atlas Kafka hook to recognize and correctly audit topic auto-creation events even when the `MetadataResponse` contains a non-zero error code.

OPSAPS-69156: Use '=' with add-opens/add-modules/add-exports

7.3.2

The JVM option syntax used by Atlas is normalized to use the equals format for module-related flags (for example, `--add-opens=...`, `--add-exports=...`, and `--add-modules=...`). This avoids failures in environments where `JAVA_TOOL_OPTIONS` does not correctly handle the space-separated variant and improves portability across JDK runtimes.

CDPD-80582: The Atlas client should retry on HTTP 500 errors

7.3.2

When the Atlas plugin attempted to bootstrap the Kafka or Schema Registry model and the Atlas server returned an HTTP 500 error (for example, error code `ATLAS-500-00-005` with the message "Failed to get the lock; another type update might be in progress"), the client did not retry the request. This caused the Atlas plugin within Schema Registry to fail initialization permanently, resulting in subsequent errors such as "schema_metadata_info: Unknown/invalid typename". The fix updates the Atlas client retry logic to also retry requests that fail with HTTP 500 (Internal Server Error), in addition to the already-handled HTTP 503 (Service Unavailable) errors.

Apache JIRA: [ATLAS-4571](#)

CDPD-69213: Export/Import, Incremental Export: When entity exported has a tag propagated from entity which is deleted, tag is not propagated to it at target

7.3.2

When creating tables with multiple levels of depth, a tag applied to the first table was propagated along the lineage. After dropping the first table and exporting the entire lineage, child tables did not have the propagated tag after import on the target cluster. The root cause was a task ordering issue in the deferred actions flow: when deferred actions were enabled during import, the add-propagation task was created before the parent entity was deleted, but ran after the deletion, finding the edge between the parent entity and the classification in a deleted state and therefore not propagating the tag to child entities. The fix adds a check in the tag propagation logic to skip the edge-status validation when import is in progress and deferred actions are enabled, ensuring that tags are correctly propagated to child entities even when the source entity has been deleted.

Apache JIRA: [ATLAS-5055](#)**CDPD-68761: Atlas 'updateTime' parameter is not updated when term is added**

7.3.2

When a glossary term was assigned to or removed from an entity, the entity's updateTime metadata field was not updated to reflect the modification. Clients or processes that relied on updateTime to detect entity changes would not recognize updates made through term assignment or removal. The fix calls the entity modification metadata update after each term assignment and dissociation operation in the glossary term processing logic, ensuring that updateTime correctly reflects the latest change.

Apache JIRA: [ATLAS-4848](#)**CDPD-67277: Improve efficiency in access request handling by optimizing classification retrieval**

7.3.2

When processing bulk classification operations on large numbers of entities, the authorization layer retrieved full classification objects for each entity to perform access checks, even though only the classification names were required. This caused unnecessary graph traversals and redundant data loading, resulting in significantly increased processing time as the number of entities and classifications grew. The fix optimizes the access request handling by using classification names directly instead of iterating over full classification objects, and by replacing the heavier entity header retrieval call that fetched all classification objects with a lighter call during authorization checks. Performance testing with 260 entities showed processing time reductions of up to 50% depending on the number of classifications involved.

CDPD-79730: Provide liveness and readiness probes in Atlas

7.3.2

Atlas now provides two REST API endpoints to support container orchestration health checks. The liveness probe endpoint (GET /api/atlas/admin/liveness) returns HTTP 200 when the Atlas service state is ACTIVE or MIGRATING, and HTTP 500 when the service requires a restart. The readiness probe endpoint (GET /api/atlas/admin/readiness) returns HTTP 200 only when the Atlas service is ACTIVE and both the index store and the backend store are active, and HTTP 500 otherwise. These endpoints enable container orchestration platforms such as Kubernetes to correctly determine when to restart an Atlas pod and when to route traffic to it.

Apache JIRA: [ATLAS-4826](#)**CDPD-70221: Atlas incremental export takes too much time**

7.3.2

When performing an incremental export of Hive database entities, Atlas traversed every connected entity in the graph and loaded the full entity object, including all extended information, for each one solely to check whether its modification timestamp fell within the export window. This resulted in data loading overhead equivalent to that of a full bootstrap export, regardless of how few entities

had actually changed. The fix introduces an optimized code path for incremental exports of Hive database entities: Atlas now reads the modification timestamp directly from the graph vertex property and loads the full entity only when the timestamp qualifies, avoiding unnecessary data loading for entities that have not been modified since the last export.

Apache JIRA: [ATLAS-4979](#)

CDPD-89927: Option to customize name of the configuration filename

7.3.2

Atlas now supports a system property named `atlas.properties` to specify the configuration filename to load, instead of being limited to `atlas-application.properties`. This enables hooks and utilities to use custom property filenames (for example, `-Datlas.properties=custom-atlas.properties`).

Apache JIRA: [ATLAS-5098](#)

Fixed Issues in Apache Avro

Review the list of Avro issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Avro issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

CDPD-88356: Migrate Apache Commons Lang dependency to version 3.18.0 to address CVE-2025-48924

7.3.2

This fix addresses a high-severity uncontrolled recursion vulnerability in Apache Commons Lang's `ClassUtils.getClass(...)` method by upgrading `commons-lang3` dependency to version 3.18.0.

CDPD-85710: Upgrade jQuery for Avro artifacts

7.3.2

Several CVEs (CVE#2020#23064, CVE#2019#5428, NSWG-ECO#328, NSWG-ECO-329) affected the jQuery version used in Avro's web artifacts. jQuery has been upgraded to a secure version address these vulnerabilities.

CDPD-82755: Restrict trusted packages in ReflectData and SpecificData

7.3.2

This fix addresses a deserialization vulnerability (CVE-2024-47561) in `ReflectData` and `SpecificData`. The change introduces the `org.apache.avro.SERIALIZABLE_PACKAGES` system property to restrict class types during serialization, protecting the host system from malicious payloads.

Apache Jira: [AVRO-3985](#)

CDPD-77463: Upgrade Avro to version 1.12.0

The Avro component has been upgraded from version 1.11.1 to version 1.12.0.

CDPD-77022: Migrate dependency to commons-lang3

7.3.2

This fix ensures that all the Avro artifacts use the centralized `commons-lang3` version.

CDPD-75119: Migrate Avro commons-io dependency to version 2.17.0

7.3.2

This fix addresses CVE-2024-47554 by upgrading `commons-io` dependency to version 2.17.0.

Fixed Issues in Cloud Connectors

Review the list of Cloud Connectors issues that are resolved in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Cloud Connectors issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.700. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

There are no fixed issues in this release.

Fixed Issues in Cruise Control

Fixed issues and resolved maintenance items for Cruise Control are addressed in Cloudera Runtime 7.3.2 and its associated service packs.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Cruise Control issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

There are no fixed issues in this release.

Fixed Issues in HDFS

Review the list of HDFS issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves HDFS issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

CDPD-75403: Fixed local information disclosure in RunJar temporary directory

7.3.2

Previously, the temporary directories created by `RunJar.run()` lacked restricted permissions by default. This could allow other local users on the shared system to access sensitive files within those directories. This issue is now resolved. The fix implements stricter access controls on the RunJar working directory, ensuring that only the owner has permission to access the contents.

Apache JIRA: [HADOOP-19031](#)

CDPD-69173: FsImage might have inaccessible nodes

7.3.2

Previously, if an FsImage was corrupted, it might have had inaccessible nodes. The FsImageValidation tool was able to identify the inaccessible nodes when validating the INodeMap. This issue is fixed now and this update enhances the FsImageValidation to remove the inaccessible nodes if the FsImage is corrupted.

Apache JIRA: [HDFS-17380](#)

CDPD-88767: OpenSSL 3.x symbols are not loading

7.3.2

Previously, Hadoop native code failed to link with OpenSSL 3.x because specific symbols—EVP_CIPHER_CTX_block_size and EVP_CIPHER_CTX_encrypting—were removed from the library's public API. This issue is now resolved. The native code has been updated to use OpenSSL 3.x compatible accessor functions, ensuring proper loading and linking.

Apache JIRA: [HADOOP-18583](#)

CDPD-78928: Backport HADOOP-16299 to avoid module access violation in LdapGroupsMapping

7.3.2

When LdapGroupsMapping was the selected group mapping method, running any commands that relies on this group mapping lead to an exception due to a java module access violation when the JVM version was above Java11. This illegal access is avoided in the group mapping implementation without functional changes.

Apache JIRA: [HADOOP-16299](#)

CDPD-84571: Backport upstream fixes for missing blocks issue

7.3.2

Backported the following upstream fixes:

Apache JIRA: [HDFS-14303](#), [HDFS-15764](#), [HDFS-16985](#), [HDFS-17342](#)

Fixed Issues in HBase

Review the list of HBase issues that are resolved in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves HBase issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

There are no fixed issues in this release.

Fixed Issues in Hive

Fixed issues for Hive are addressed in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Hive issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

CDPD-97246: Avro schema literal logging at INFO level

7.3.2

Previously, the Avro deserializer logged the schema literal at the INFO level.

This issue is now resolved by changing the log level to DEBUG.

Apache Jira: [HIVE-22606](#)

CDPD-96649: Incorrect aggregate statistics when direct SQL batch retrieval is enabled

7.3.2

Previously, when hive.metastore.direct.sql.batch.size config was greater than 0, the system failed to merge column statistics correctly if the number of partitions or columns exceeded that batch size.

This issue is now resolved by ensuring the statistics are properly merged during the retrieval process, preventing redundant entries.

Apache Jira: [HIVE-29203](#)

CDPD-95834: Hive Metastore backend database schema out of date

7.3.2

Previously, the Hive Metastore (HMS) backend database schema files required updates for compatibility with the latest version.

This issue is now resolved by updating the schema files to version 7.3.2.0 to ensure a successful upgrade.

CDPD-95681: Hive Beeline connection failure due to SSL certificate hostname mismatch

7.3.2

Previously, the Beeline client failed to establish a connection to HiveServer2 in ZooKeeper High Availability (HA) setups.

This issue is now resolved by ensuring that the SSL certificates correctly match the DNS hostnames, allowing the secure connection to be verified and established.

CDPD-93766: Missing results during anti-join conversion

7.3.2

Previously, queries involving specific join conditions returned empty results instead of the expected data when anti-join conversion was enabled.

This issue is now resolved by preventing anti-join conversion in these specific scenarios to ensure query result accuracy.

Apache Jira: [HIVE-29175](#)

CDPD-93756: SHOW COMPACTIONS output filtering

7.3.2

Previously, the SHOW COMPACTIONS output included all historical information, which caused the display to become unwieldy when many partitions or history lines existed.

This issue is resolved by adding the ability to filter the command output by database, table, partition, and compaction type or state.

Apache Jira: [HIVE-13353](#)

CDPD-93617: Duplicate records during minor compaction

7.3.2

Previously, when the Hive Metastore (HMS) crashed, active compaction jobs were incorrectly reset.

This issue is now fixed by updating the compactor cleaner to address duplicate directories.

Apache Jira: [HIVE-29210](#)

CDPD-93432: Incorrect results for anti-join queries

7.3.2

Previously, the HiveAntiJoin rule incorrectly replaced IS NULL filters on nullable columns, which resulted in missing records in the query output.

This issue is now fixed by improving the logic within the HiveAntiSemiJoinRule to ensure accurate query plans.

Apache Jira: [HIVE-29176](#)

CDPD-93431: Incorrect results for n-way joins

7.3.2

Previously, queries produced incorrect results when an n-way join contained a combination of both anti and outer joins.

This issue is now resolved by extending the CommonJoinOperator to properly support anti joins when they are used alongside outer joins in n-way join operations.

Apache Jira: [HIVE-29290](#)

CDPD-93428: Manifest files appearing in read queries

7.3.2

Previously, temporary directories containing direct insert manifest files used a prefix that allowed them to be included in concurrent read queries.

This issue is now resolved by hiding the direct insert manifest directory from read queries, ensuring only valid data files are processed.

Apache Jira: [HIVE-29297](#)

CDPD-93425: Data loss during minor compaction

7.3.2

Previously, query-based minor compaction incorrectly used the minimum open write ID as a lower bound for selecting data files.

This issue is now fixed by removing the incorrect check against the minimum open write ID, allowing the high-watermark to correctly define the compaction range.

Apache Jira: [HIVE-29272](#)

CDPD-92923 / CDPD-92586: Memory leak in Hive Metastore REST Catalog

7.3.2

Previously, the Hive Metastore (HMS) REST Catalog leaked memory during `ALTER TABLE` authorization checks. Each check created new catalog instances and handler pools with JMX enabled, which prevented the system from reclaiming memory.

This issue is now fixed by reusing existing catalog instances and disabling JMX for the handler pool to allow the system to recover resources.

CDPD-92499: Hive LDAP authentication and ZooKeeper connections

7.3.2

Previously, Hive Lightweight Directory Access Protocol (LDAP) authentication failed when connecting to a SASL-enforced ZooKeeper instance.

This issue is now resolved by updating the service components to support simultaneous authentication configurations and ensuring correct credential handling during service discovery.

Apache Jira: [HIVE-29138](#)

CDPD-92478: Timestamp processing in MetaStoreUtils

7.3.2

Previously, Hive Metastore utilities used local time zone settings to convert between timestamps and strings.

This issue is now fixed by using UTC time zone and the `java.time.Instant` class to process timestamps, which ensures that time points are represented accurately regardless of local time zone rules.

Apache Jira: [HIVE-28337](#)

CDPD-91415: WebHCat and Python script compatibility with Python 3

7.3.2

Previously, `hcat.py` and various Python scripts used in `q` files contained syntax that was incompatible with Python 3.

This issue is now resolved by updating the Python scripts to use Python 3-compatible syntax.

Apache Jira: [HIVE-25817](#)

CDPD-90670: Incorrect results for queries with multiple lateral view operations

7.3.2

Previously, queries that used two or more `LATERAL VIEW` explode operations along with a `WHERE` clause returned incorrect results when Cost-Based Optimization (CBO) was enabled.

This issue is now resolved by updating the logic to correctly identify separate table aliases for lateral view columns, ensuring that filters are applied accurately.

Apache Jira: [HIVE-29084](#)

CDPD-90303: Incorrect results from a CASE expression

7.3.2

A query that used a `CASE` expression to conditionally return values produced an incorrect result. The query plan incorrectly folded the `CASE statement` into a `COALESCE` function, which led to a logic error that filtered out some of the expected results.

This issue is addressed by adding a more strict check when converting `CASE` expressions into `COALESCE` during query optimization.

Apache Jira: [HIVE-24902](#)

CDPD-89462: Performance degradation for wide tables in DirectSqlUpdatePart

7.3.2

Previously, updating or inserting partition statistics for tables with a high number of columns and partitions was slow.

This issue is now resolved by improving the hashing logic to ensure faster data retrieval and insertion, even for tables with thousands of columns and partitions.

Apache Jira: [HIVE-29165](#)

CDPD-88987: Improved performance for adding columns to partitioned tables

7.3.2

Previously, metadata operations for adding columns to tables with a high number of partitions and columns were slow because the system utilized a less optimized implementation for partition updates.

This issue is addressed by implementing a more efficient batch processing method for partition updates, which utilizes optimized metadata queries to improve performance for tables with many partitions.

Apache Jira: [HIVE-28956](#)

CDPD-88981: Performance degradation during column addition with cascade

7.3.2

Previously, adding columns to a table using the `CASCADE` command resulted in slower performance after optimizations for metadata storage were enabled.

This issue is now resolved. The fix includes an optimized method to reuse column descriptors across partitions, which restores performance levels during the column addition process.

Apache Jira: [HIVE-29042](#)

CDPD-88166: Query failure during JDBC filter optimization

7.3.2

Previously, certain queries failed with a class-cast error during the optimization phase.

This issue is now resolved by updating the query optimization rules to ensure that relational operators are correctly identified and processed.

Apache Jira: [HIVE-25356](#)

CDPD-87266: Query failure during Tez execution

7.3.2

Previously, Hive queries failed during execution after an upgrade. This resulted in a vertex failure and prevented queries from completing successfully.

This issue is now resolved by updating the execution engine to correctly instantiate internal query split generators during the initialization process.

CDPD-84149: MariaDB connector recognition failure

7.3.2

Previously, Hive failed to recognize the MariaDB connector even when the driver was present.

This issue is now fixed.

CDPD-83461: Query failure when using stack function with union operations

7.3.2

Previously, queries utilizing the `STACK` function in combination with `UNION` operations failed with an internal error during the compilation phase.

This issue is now resolved by updating the query optimization logic to correctly handle the stack function during union operations, preventing the internal processing error.

Apache Jira: [HIVE-29029](#)

CDPD-83334: Improving performance for alter partition operations

7.3.2

When altering partitions, the system used Java Data Objects (JDO) updates, which required fetching all fields of old partitions and produced redundant queries, leading to slower performance.

This issue is resolved by implementing direct SQL for altering partitions.

Apache Jira: [HIVE-27530](#)

Direct SQL failure during partition alterations

7.3.2

Previously, direct SQL for partition alterations failed in certain database environments due to Character Large Object (CLOB) casting errors and missing boolean type conversion checks.

This issue is resolved by updating the direct SQL logic to handle CLOB types correctly and ensuring proper boolean type conversions during batch updates.

Apache Jira: [HIVE-28271](#)

CDPD-80146: Group by alias query failures

7.3.2

Previously, the `hive.runtime.dialect.enable` property was enabled by default, which caused the `hive.groupby.position.alias` property to be ignored.

This issue is resolved by setting the `hive.runtime.dialect.enable` property to false by default. Hive now correctly respects the `hive.groupby.position.alias` configuration.

CDPD-79144: Incorrect schema version in Hive schema initialization script for MySQL

7.3.2

Cluster creation with 7.3.1 fails due to an incorrect database schema in the `CDH_VERSION` table.

The issue was addressed by correcting the schema version in the Hive schema initialization script, ensuring successful cluster creation.

CDPD-78337: Merge task not invoked for external CTAS queries on object stores

7.3.2

Previously, the merge task was not invoked for external Create Table As Select (CTAS) queries when using S3 or other object stores.

This issue is now resolved by ensuring the merge task is correctly invoked after optimization for external CTAS queries.

Apache Jira: [HIVE-27536](#)

CDPD-78329: HiveServer2 runs out of memory with multiple parallel queries with fetch task

7.3.2

Previously, HiveServer2 may run out of memory when multiple parallel queries use fetch task caching (`hive.fetch.task.caching=true`). This causes queries to fail and HiveServer2 to crash.

The issue was addressed by reducing the default value of `hive.fetch.task.conversion.threshold` from one GB to 200MB, preventing excessive memory usage and improving stability.

CDPD-77869: Iceberg table data written to HDFS instead of S3 in RAZ-enabled clusters

7.3.2

Previously, when you configured a cluster with Ranger Remote Authorization Board (RAZ) and updated configurations to use S3, data for Iceberg tables was unexpectedly written to the HDFS external table location. This occurred because Iceberg tables, which are treated as external tables, defaulted to the database LOCATION property that still pointed to HDFS if the database was created prior to the S3 switch.

This issue is now fixed by ensuring that external tables correctly align with the intended S3 storage paths. You can now update the database location to point to the S3 bucket to ensure all future external tables default to the cloud storage.

CDPD-75665: Importing a table generates a DDL with an incorrect location

7.3.2

When creating a table using the IMPORT command, this table's partitions could point to an incorrect location, that is, an external imported table can have his partitions located under the managed warehouse directory, which violates the `metastore.warehouse.external.dir` and `metastore.warehouse.dir` that intend to host different type of tables.

Apache Jira: [HIVE-28580](#)

Hive query execution failure due to AM container exit on lost node with Exit code -100

7.3.2

Hive query failed when ApplicationMaster container was lost

Previously, when running a Hive query, a failed ApplicationMaster (AM) container did not trigger a DAG retry and caused the query execution to fail if the failure message included diagnostic information with a line break.

This issue is now resolved by automatically re-executing the DAG if the AM fails.

Apache Jira: [HIVE-28093](#)

CDPD-74539: MariaDB falls back to MySQL in Hive

7.3.2

Hive downstream had errors in supporting MariaDB.

The issue was addressed by making MariaDB automatically fall back to MySQL.

CDPD-66731: Hive Metastore query failure during Zero Downtime Upgrade

7.3.2

Previously, during a Zero Downtime Upgrade (ZDU) from version 7.2.17 to 7.2.18, long-running queries such as `INSERT INTO` statements failed with a `MetaException`.

This issue is now fixed by ensuring that transaction blocks are correctly managed during the statistics update task, preventing the "current transaction is aborted" error.

CDPD-60770: Passwords with special characters fail to connect with Beeline

7.3.2

When you used a password containing special characters like `#`, `^`, or `;` in a JDBC URL for a Beeline connection, the connection failed with a 401 error. This happened because Beeline did not correctly interpret these special characters in the password.

This issue is resolved by introducing a new method to reparse the password from the original JDBC URL, allowing Beeline to correctly handle and authenticate passwords containing special characters.

Apache Jira: [HIVE-28805](#)

CDPD-58428: Bucket Map Join hangs when source vertex parallelism changes

7.3.2

Previously, a Bucket Map Join could hang if the parallelism of a source vertex was modified by automated reducer parallelism.

This issue is now fixed by disabling automated reducer parallelism for vertices that serve as a source for a Bucket Map Join.

Apache Jira: [HIVE-27078](#)

CDPD-58428: Incorrect results in map-side Sort-Merge Bucket Join with different bucket sizes

7.3.2

Previously, map-side Sort-Merge Bucket (SMB) Joins returned incorrect results when joining two tables with different bucket counts (for example, joining a table with two buckets to a table with three buckets).

This issue is now fixed by implementing a new routing algorithm that ensures bucket N of a small table is correctly mapped to bucket M of a large table based on the greatest common divisor of their bucket sizes.

Apache Jira: [HIVE-27357](#)

CDPD-50060: Configurable filter for partition metadata properties in Hive Metastore

7.3.2

Previously, Hive Metastore (HMS) API calls failed with a `TTransportException` (`MaxMessageSize` reached) when processing tables with large partition metadata.

This issue is now fixed by providing a configurable filter that excludes unnecessary properties from `listPartitions` API responses. This change reduces the metadata payload size, prevents connection timeouts, and improves the performance of metadata operations.

Apache Jira: [HIVE-27114](#)

CDPD-44551: Avro table import or download fails with ODBC driver due to missing property

7.3.2

The absence of `metastore.storage.schema.reader.impl` caused Avro table import or download failures in Cloudera runtime 7.1.7 when using the ODBC driver.

The issue was addressed by setting `metastore.storage.schema.reader.impl` to `org.apache.hadoop.hive.metastore.SerdeStorageSchemaReader` by default.

Apache Jira: [HIVE-26952](#)

CDPD-92208: Query failure when selecting data from views with bracketed definitions

7.3.2

Previously, a `SELECT` query against a view failed with a `SemanticException` if the view was created with specific column names and a definition enclosed in brackets.

This issue is now fixed by ensuring that the compiler does not add extra brackets if the view definition is already enclosed.

Apache Jira: [HIVE-26493](#)

CDPD-83530: Task commits were allowed despite an exception being thrown in the Tez processor

7.3.2

A communication failure between the coordinator and executor caused a running task to terminate, resulting in a `java.lang.InterruptedException` being thrown by the `ReduceRecordProcessor.init()`. Despite this exception, the process still allowed the task to be committed and generated a commit manifest.

This issue has now been resolved. The fix ensures that outputs are not committed if an exception is thrown in the Tez processor.

Apache Jira: [HIVE-28962](#)

CDPD-89414: Incorrect results for window functions with IGNORE NULLS

7.3.2

When you used the `FIRST_VALUE` and `LAST_VALUE` window functions with the `IGNORE NULLS` clause while vectorization was enabled, the results were incorrect. This occurred because the vectorized execution engine did not properly handle the `IGNORE NULLS` setting for these functions.

This issue is addressed by modifying the vectorized processing for `FIRST_VALUE` and `LAST_VALUE` to correctly respect the `IGNORE NULLS` clause, ensuring the same results are produced whether vectorization is enabled or disabled.

Apache Jira: [HIVE-29122](#)

CDPD-85600: Select queries with ORDER BY fail due to compression error

7.3.2

When you ran a Hive `SELECT` query with an `ORDER BY` clause, it failed with a `java.io.IOException` and `java.lang.UnsatisfiedLinkError` related to the `zlib` decompressor.

The issue was addressed by ensuring the `zlib` native library is correctly loaded.

Apache Jira: [HIVE-28805](#)

CDPD-90301: Stack overflow error from queries with OR and MIN filters

7.3.2

Queries, cause a stack overflow error when they contained multiple `OR` conditions on the same expression, such as `MINUTE(date_) = 2 OR MINUTE(date_) = 10`.

This issue is addressed by modifying the `HivePointLookupOptimizerRule` to keep the original order of expressions and to check if a merge can be performed before creating a new expression.

Apache Jira: [HIVE-29208](#)

DWX-20754: Invalid column reference in lateral view queries

7.3.2

The virtual column `BLOCK__OFFSET__INSIDE__FILE` fails to be correctly referenced in queries using lateral views, resulting in the error:

```
FAILED: SemanticException Line 0:-1 Invalid column reference 'BLOCK__OFFSET__INSIDE__FILE'.
```

This issue is now resolved.

Apache Jira:[HIVE-28938](#)

Fixed Issues in Cloudera Data Explorer (Hue)

Fixed issues for Cloudera Data Explorer (Hue) are addressed in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Impala issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

CDPD-73189: Not all tables are loading in the left assist panel

7.3.2

Previously, in the Data Explorer left assist panel, only the initial 5,000 tables were loaded because the default value of `max_catalog_sql_entries` was set to 5000. Increasing this configuration in the Cloudera Data Explorer (Hue) Safety Valve only allowed loading up to 10,000 tables due to the `max_rows` parameter being hardcoded to 10000. This issue is now resolved by fixing the hardcoded limitation. You can now load more than 10,000 tables in the left assist panel.

CDPD-71098: Complex Hive or Impala data type sampling fails in Data Explorer

7.3.2

Previously, sampling complex Hive or Impala data types in Data Explorer did not work correctly. This issue is now fixed. The resolution enables accurate sampling and autocomplete functionality for nested complex types, which improves the query experience.

CDPD-48151: Manual table creation using the importer

7.3.2

Previously, a JavaScript error occurred when you used the Manually option in the importer to create a table. This issue prevented fields from being displayed. This issue is now fixed. To create a table without importing data, you can also use the SQL editor.

CDPD-47230: Missing SQL information in Table Browser for Hive views

7.3.2

Previously, the `viewSQL` feature in the Table Browser displayed an empty window for views created by using Hive. This issue occurred when the source type was set to Hive, while the feature worked correctly for Impala. This issue is now fixed. You can now view the SQL information for Hive views in the Table Browser as expected.

CDPD-92829: Job Browser error and traceback exposure

7.3.2

Previously, Job Browser displayed an error if the `POST` format parameter contained empty or invalid values. This issue caused a JSON serialization failure and exposed sensitive system information. This issue is now fixed. The resolution includes validated parameter inputs and improved error handling to prevent the exposure of sensitive system information.

CDPD-77831: Workflow log icon fails to redirect in Data Explorer

7.3.2

Previously, the log icon in the Task tab for workflows did not redirect to the logs. When you clicked the icon during a workflow run, the UI remained on the same page without displaying an error, although a `TypeError` related to `lodash.trimRight` occurred in the background. This issue is now fixed. You can now access workflow logs through the UI as expected.

CDPD-74389: File and folder creation fails in Data Explorer S3 file browser

7.3.2

Previously, you could not create files or folders in the Data Explorer S3 File Browser on the AWS-Fedramp-RAZ environment. These actions failed to complete without displaying an error message in the UI. This issue is now fixed.

CDPD-90906: Database connection errors in metrics module

7.3.2

Previously, the metrics module encountered the `cx_Oracle.DatabaseError: DPI-1010: disconnect error` when the Cloudera Manager metrics API queried active users. This issue caused database connections to be lost. This issue is now fixed. The resolution improves connection stability and robustness against database disconnect errors.

CDPD-89478: Oozie UI displays start and end times incorrectly after upgrade

7.3.2

Previously, the Oozie UI intermittently displayed the start time and end time fields in reverse order. Additionally, the Run button remained inactive until you clicked Save, even if you made no changes. This issue is resolved by correcting the field display order and ensuring that the Run button is active for schedules without requiring an unnecessary save action.

CDPD-58142: Query pre-population fail in the SQL editor

7.3.2

Previously, when you clicked Re-Execute to rerun a query from the Job Browser Queries Query Details page, the query did not appear in the SQL editor as expected. This issue occurred because the editor loaded multiple times, overwriting the query text. Additionally, the page change was not reflected in the URL. This issue is now fixed. The SQL editor now correctly populates and retains the query text when you navigate from the Job Browser.

CDPD-65749: Table Browser navigation error with Knox

7.3.2

Previously, clicking the back button when navigating the Table Browser directory structure through Knox resulted in a 404 Error. This issue is now fixed. The resolution improves navigation stability and ensures the UI correctly displays the previous page.

CDPD-81025: Performance issues when querying large datasets with Oracle or MySQL queries

7.3.2

Previously, when querying large datasets using Oracle or MySQL backends, the `QueryProcessor` failed to correctly apply the SQL limit clause. This caused the system to fetch all data records instead of the restricted set required for pagination. As a result, the user interface experienced significant loading delays and performance issues when displaying query results. This issue is now resolved.

Fixed Issues in Impala

Fixed issues for Impala are addressed in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Impala issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

CDPD-98207: Impala crashing on the Web UI for failed queries

7.3.2

Previously, Impala crashed when you accessed the query summary or JSON plan through the Web UI for queries that failed before execution. This occurred during scenarios such as a Create Table As Select (CTAS) failure or when admission control rejected a query.

This issue is addressed by ensuring the system correctly handles missing execution summaries. This issue is now fixed.

Apache Jira: [IMPALA-14791](#)

CDPD-97786: Excessive partition events during table-level operations

7.3.2

Previously, certain table-level operations, such as dropping incremental statistics or setting/unsetting cached properties, triggered an individual ALTER PARTITION event for every partition in a table.

This issue is addressed by implementing bulk updates for partitions.

Apache Jira: [IMPALA-13599](#)

CDPD-97187: Deprecation warnings in impala-shell with Python 3.11

7.3.2

Previously, when running `impala-shell` by using Python 3.11 or newer, the command output displayed DeprecationWarning messages related to `ssl.PROTOCOL_TLS` and `ssl.match_hostname()`. These warnings were triggered by underlying library dependencies.

This issue is now resolved by updating the handling of SSL protocols and validating logic to be compatible with newer Python versions, which eliminates these warning messages from the shell output.

Apache Jira: [IMPALA-12219](#)

CDPD-91994: Stale query IDs in catalog logs

7.3.2

Previously, catalog logs for `getPartialCatalogObject` certain metadata requests displayed incorrect query IDs.

This issue is addressed by ensuring that each request is associated with its correct query ID. The system now automatically clears the identification after the request finishes to prevent stale information from appearing in later logs.

Apache Jira: [IMPALA-14494](#)

CDPD-82673: RSASSA-PSS certificate signature schema is now supported for server certificates

7.3.2

Previously, if you used a certificate with the RSASSA-PSS signature algorithm for kRPC communication, the connection failed.

The fix includes an updated OpenSSL function that correctly identifies the hash algorithm for RSASSA-PSS certificates.

Apache Jira: [IMPALA-14038](#)

CDPD-89852: Crash when casting timestamp strings with timezone offsets to DATE

7.3.2

Attempting to cast a timestamp string that included a timezone offset (like "+08:00" in "2025-08-31 06:23:24.9392129 +08:00") to the DATE data type would cause a crash.

This issue is addressed by adding a check to ensure that the timestamp string length does not exceed the maximum length of the default date-time format. Longer strings will now use a lazily-created format, which prevents the crash.

Apache Jira: [IMPALA-14383](#)

CDPD-89730: Impala daemon crashed during scans with high logging levels

7.3.2

Previously, the Impala daemon experienced a null pointer dereference in the BaseSequenceScanner component when the logging level was set to 2 or higher, leading to crashes in release builds.

This issue is resolved by correcting the pointer handling in the sequence scanner to ensure safe memory access when high-level logging is active.

Apache Jira: [IMPALA-14382](#)

CDPD-89346: Enhanced join strategy selection for large clusters

7.3.2

The query planner's cost model for broadcast joins can be skewed by the number of nodes in a cluster. This led to suboptimal join strategy choices, especially in large clusters with skewed data where a partitioned join was chosen over a more efficient broadcast join.

This issue is now resolved by introducing the `broadcast_cost_scale_factor` query option as an additional tuning option besides query hint to override query planner decision.

Apache Jira: [IMPALA-14263](#)

CDPD-89132: Tables incorrectly dropped by stale HMS events after global metadata invalidation

7.3.2

Previously, a stale event such as DropTable or AlterTableRename post global `INVALIDATE METADATA` command could cause tables to be unintentionally dropped.

This issue is resolved by tracking the `createEventId` as the current HMS event ID for all tables during a global reset.

Apache Jira: [IMPALA-14330](#)

CDPD-79111: Authentication failure in impala-shell with 76 character LDAP passwords

7.3.2

Previously, when you used `impala-shell` with the HS2-HTTP protocol and a 76 character LDAP password, the connection failed with a value error.

This issue is resolved by an updated encoding method that handles long password strings without inserting line breaks, ensuring that the authorization header remains valid for the server.

Apache Jira: [IMPALA-13746](#)

CDPD-92001: Metadata loading performed sequentially in local catalog mode

7.3.2

Previously, when a query accessed multiple unloaded tables in local catalog mode, Impala triggered metadata loading for those tables sequentially.

This issue is resolved by parallelizing table loading during query compilation. A new startup flag, `max_stmt_metadata_loader_threads`, is introduced to control the number of threads used for loading metadata, with a default value of 8 threads per query. If only one table requires loading or if the thread pool is unavailable, the system automatically falls back to sequential loading.

[IMPALA-14447](#)**CDPD-79241: Incorrect query results for Iceberg V2 tables**

7.3.2

Previously, when you ran complex queries involving multiple subqueries on Iceberg V2 tables, the system sometimes returned incorrect results.

This issue is now resolved. The fix includes a new internal mechanism to track and apply count optimizations.

Cookie-Based authentication support for JWT tokens

7.3.2

When JWT tokens are used for authentication, every HTTP request within a session requires token verification. If these tokens have a short lifespan, it can lead to authentication failures and disrupt session continuity.

This issue is now resolved by using authentication cookies, which generally have a longer lifespan (configured through the `max_cookie_lifetime_s` flagfile option) and can remain valid for the duration of the session. This enables subsequent authentication requests to rely on cookies rather than repeatedly verifying the JWT token.

Apache Jira: [IMPALA-13813](#)

CDPD-80798: Stable Catalogd initialization in HA mode

7.3.2

Catalogd initialization previously might timeout to complete in high availability mode. This happened because metadata operations started prematurely, blocking Catalogd from becoming active.

This issue is resolved by ensuring Catalogd determines HA state before starting metadata operations in HA mode. This prevents blocking issues and ensures a stable startup.

Apache Jira: [IMPALA-13850](#)

CDPD-83059: Optimized Impala Catalog cache warmup

7.3.2

Impala's Catalogd previously started with an empty cache. This led to slow query startup for important tables and affected high availability failovers.

This issue is resolved by adding new settings to pre-load specific tables into the Catalogd cache in the background. This ensures faster query startup and smoother high availability failovers.

Apache Jira: [IMPALA-14074](#)

CDPD-87222: Consistent TRUNCATE operations for external tables

7.3.2

Impala's TRUNCATE operations on external tables previously did not consistently delete files in subdirectories, even when recursive listing was enabled.

This issue is resolved by ensuring Impala uses the HMS API for TRUNCATE operations by default.

Apache Impala: [IMPALA-14189](#), [IMPALA-14224](#)

DWX-21855: Impala Executors fail to gracefully shutdown

7.3.2

During graceful shutdown Impala executors wait for running queries to finish up to the graceful shutdown deadline (`--shutdown_deadline_s`). During graceful shutdown the istio-proxy container on Impala executor pod was getting terminated immediately and as a result the executors were not reachable and were removed from the Impala cluster membership resulting in cancellation of running queries.

This issue is now resolved by making sure istio-proxy container's lifecycle doesn't impact executor's cluster membership.

IMPALA-14263: Enhanced join strategy for large clusters

7.3.2

The query planner's cost model for broadcast joins can be skewed by the number of nodes in a cluster. This could lead to suboptimal join strategy choices, especially in large clusters with skewed data where a partitioned join was chosen over a more efficient broadcast join.

This issue is now resolved by introducing the `broadcast_cost_scale_factor` query option as an additional tuning option besides query hint to override query planner decision. To set it cluster-wide for all queries, add the following key-value to the `default_query_options` startup option:

```
broadcast_cost_scale_factor=<less than 1.0>
```

Apache Jira: [IMPALA-14263](#)

IMPALA-11402: Fetching metadata for tables with huge numbers of files no longer fails with OutOfMemoryError

7.3.2

Previously, when Impala Coordinator tried to fetch file metadata for extremely large tables (those with millions of files or partitions), the Impala Catalog service would attempt to return all the file details at once. This often exceeded the Java memory limits, causing the service to crash with an `OutOfMemoryError`.

This issue is addressed by configuring the Catalog service to limit the number of file descriptors included in a single `getPartialCatalogObject` response. A new configuration flag, `catalog_partial_fetch_max_files`, is introduced to define the maximum number of file descriptors allowed per response (with a default of 1,000,000 files).

If a request exceeds this limit, the Catalog service will truncate the response and return metadata for only a subset of the requested partitions. The coordinator is now designed to detect this truncated response and automatically send new batch requests to fetch the remaining partitions until all required metadata is retrieved. This change ensures that the coordinator can successfully fetch and process the metadata for extremely large tables without crashing due to memory limits.

Apache Jira: [IMPALA-11402](#)

CDPD-77261: Impala can now read Parquet integer data as DECIMAL after schema changes

7.3.2

Previously, if you changed a column type from an integer (`INT` or `BIGINT`) to a `DECIMAL` using `ALTER TABLE`, Impala could fail to read the original Parquet data files. This happened because the files lacked the specific metadata (logical types) Impala expected for decimals, resulting in an error.

Impala is now more flexible when reading Parquet files following schema evolution. If Impala encounters an integer type but the schema expects a `DECIMAL`, it automatically assumes a suitable decimal precision and scale, allowing you to successfully query the updated table:

- `INT32` is read as `DECIMAL(9, 0)`.
- `INT64` is read as `DECIMAL(18, 0)`.

This change supports common schema evolution practices by allowing you to update column types without manually rewriting old data files.

Apache Jira: [IMPALA-13625](#)

IMPALA-12927: Impala can now correctly read BINARY columns in JSON tables

7.3.2

Previously, Impala couldn't correctly read BINARY columns in JSON tables, often resulting in errors or incorrect data. This happened because Impala assumed the data was always Base64 encoded, which wasn't true for files written by older Hive versions.

Impala now supports a new table property, 'json.binary.format' (BASE64 or RAWSTRING), and a query option, JSON_BINARY_FORMAT, to explicitly define the binary encoding. This ensures Impala reads the data correctly. If no format is specified, Impala will now return an error instead of risking silent data corruption.

JIRA Issue: [IMPALA-12927](#)

Fixed Issues in Apache Iceberg

Review the list of Iceberg issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Iceberg issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

CDPD-78686: Iceberg tables created in 7.2.17 are not captured in Atlas when using 7.2.18 or 7.3.1 Atlas server

7.3.2

This issue occurred due to incompatibility between Data Services, Hive, and Impala hooks in 7.2.17 and the Atlas server in 7.2.18 and 7.3.1.

This fix resolves the compatibility issue and Iceberg tables created in 7.2.17 are now correctly captured and displayed in the Atlas UI in later versions.

CDPD-97171: Concurrency issues between compaction and concurrent write operations

7.3.2

This fix resolves an issue where compaction conflicted with concurrent write operations causing data corruption, by improving concurrency handling to ensure stable operations and data consistency.

Apache Jira: [HIVE-29437](#)

CDPD-74040: Dropping Iceberg table with complex type containing timestamp fails

7.3.2

This fix resolves an issue where dropping an Iceberg table with complex types such as array, map, or struct containing timestamp fields failed due to unsupported Hive timestamp type handling, by ensuring proper type conversion and successful table deletion.

CDPD-89402: Event processing invalidates Iceberg tables due to reload failures

7.3.2

This fix resolves an issue where CatalogServiceCatalog.reloadTableIfExists() resulted in a ClassCastException during event processing, invalidating Iceberg tables and triggering full table reloads instead of incremental loading.

Apache Jira: [IMPALA-14358](#)

CDPD-72383: COUNT(*) optimization returns incorrect results in UNION queries on Iceberg V2 tables

7.3.2

A faulty COUNT(*) query optimization caused incorrect results in UNION queries on Iceberg V2 tables that have delete files, leading to inconsistent query outputs.

The fix corrects the optimization logic to ensure accurate results across UNION queries, including scenarios involving post data changes.

Apache Jira: [IMPALA-13756](#) [IMPALA-13249](#)

CDPD-81076: LEFT ANTI JOIN fails on Iceberg V2 tables with Delete files

7.3.2

Queries using a LEFT ANTI JOIN fail with an AnalysisException if the right-side table is an Iceberg V2 table containing delete files. For example, consider the following query:

```
SELECT * FROM table_a a
LEFT ANTI JOIN iceberg_v2_table b
ON a.id = b.id;
```

The error Illegal column/field reference'b.input_file_name' of semi-/anti-joined table 'b' is displayed because semi-joined tuples need to be explicitly made visible for paths pointing inside them to be resolvable.

The fix updates the IcebergScanPlanner to ensure that the tuple containing the virtual fields is made visible when it is semi-joined.

Apache Jira: [IMPALA-13888](#)

CDPD-78427: Enable MERGE statement for Iceberg tables with equality deletes

7.3.2

This patch fixes an issue that caused MERGE statements to fail on Iceberg tables that use equality deletes.

The failure occurred because the delete expression calculation was missing the data sequence number, even though the underlying data description included it. This mismatch caused row evaluation to fail.

The fix ensures the data sequence number is correctly included in the result expressions, allowing MERGE operations to complete successfully on these tables.

Apache Jira: [IMPALA-13674](#)

CDPD-77773: Tolerate missing data files during Iceberg table loading

7.3.2

This fix addresses an issue where an Iceberg table would fail to load completely if any of its data files were missing from the file system. This TableLoadingException left the table in an incomplete state, blocking all operations on it.

Impala now tolerates missing data files during the table loading process. An exception will only be thrown if a query subsequently attempts to read one of the specific files that is missing.

This change allows other operations that do not depend on the missing data—such as ROLLBACK, DROP PARTITION, or SELECT statements on valid partitions—to execute successfully.

Apache Jira: [IMPALA-13654](#)

CDPD-78508: Skip reloading Iceberg tables when metadata JSON file is the same

7.3.2

This patch optimizes metadata handling for Iceberg tables, particularly those that are updated frequently.

Previously, if an event processor was lagging, Impala might receive numerous update events for the same table (for example, 100 events). Impala would attempt to reload the table 100 times, even if the table's state was already up-to-date after processing the first event.

With this fix, Impala now compares the path of the incoming metadata JSON file with the one that is currently loaded. If the metadata file location is the same, Impala skips the reload, correctly assuming the table is already unchanged. This significantly reduces unnecessary metadata processing.

Apache Jira: [IMPALA-13718](#)

CDPD-82415: TABLESAMPLE clause of the COMPUTE STATS statement has no effect on Iceberg tables

7.3.2

This fix resolves a regression introduced by [IMPALA-13737](#). For example, the following query scans the entire Iceberg table to calculate statistics, whereas it should ideally use only about 10% of the data.

```
COMPUTE STATS t TABLESAMPLE SYSTEM system(10);
```

This fix introduces proper table sampling logic for Iceberg tables, which can be utilized for COMPUTE STATS. The sampling algorithm previously located in `IcebergScanNode.GetFilesSample()` is now relocated to `FeIcebergTable.Utils.GetFilesSample()`.

Apache Jira: [IMPALA-14014](#)

CDPD-85228: IllegalStateException with Iceberg table with DELETE

7.3.2

Running a query on an Iceberg table fails with an `IllegalStateException` error in the following scenario:

- The Iceberg table has delete files for every data file (no data files without delete files) AND
- An anti-join operation is performed on the result of the Iceberg delete operation (IcebergDeleteNode or HashJoinNode)

This fix resolves the issue by setting the `TableRefIds` of the node corresponding to the Iceberg delete operation (IcebergDeleteNode or HashJoinNode) to only the table reference associated with the data files, excluding the delete files.

Apache Jira: [IMPALA-14154](#)

CDPD-87405: Error unnesting arrays in Iceberg tables with DELETE files

7.3.2

The following error occurred when unnesting a nested array (a 2D array) from an Iceberg table. This issue was triggered specifically when the table contained delete files for some, but not all, of its data files.

```
Filtering an unnested collection that comes from a UNION [ALL] is not supported yet.
```

Reading an Iceberg table with this mixed data and delete file configuration creates a UNION ALL node in the query execution plan. The system had a check that explicitly blocked any filtering on an unnested array.

This fix relaxes the validation check, allowing the operation to proceed if all UNION operands share the same tuple IDs. This ensures the query can successfully unnest the array.

Apache Jira: [IMPALA-14185](#)

Fixed Issues in Kafka

Fixed issues and resolved maintenance items for Kafka are addressed in Cloudera Runtime 7.3.2 and its associated service packs.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Kafka issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

OPSAPS-71679: Ranger resource-based policy creation attempted on each service restart

7.3.2

On each service restart, Kafka, Kafka Connect, and Schema Registry attempted to create Ranger resource-based policies even when the Ranger policy cache was already present. Resource-based policy creation is now skipped when the policy cache file exists, reducing unnecessary Ranger operations on restart.

OPSAPS-75094: Broker rolling restart checks fail on FIPS-enabled clusters with JDK 11 or higher

7.3.2

On FIPS-enabled clusters with JDK 11 or higher, CSD_JAVA_OPTS was not propagated to the Kafka broker rolling restart check scripts, which caused the checks to fail. The rolling restart check scripts now receive the CSD_JAVA_OPTS value on FIPS-enabled clusters.

DBZ-4990: The Debezium Db2 Source connector does not support schema evolution

7.3.2

This issue is fixed. For more information, see [DBZ-4990](#).

None.

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

7.3.2

Rolling restart checks for Kafka Connect failed when mutual TLS was required (ssl.client.auth set to required), even though the service was healthy. This issue is fixed.

Fixed Issues in Kudu

Fixed issues for Kudu are addressed in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Kudu issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

CDPD-97140: Range partition recreation for specific table schemas

7.3.2

Previously, dropping and immediately adding back the same range partition failed for certain table schemas. This issue occurred because the system catalog did not correctly normalize range-specific hash schemas when they matched the table-wide hash schema. This lack of normalization caused inconsistencies in how the partition information was stored and interpreted. This issue is now fixed. Kudu now correctly normalizes and stores range partition specifications, ensuring that partitions can be recreated as expected.

CDPD-95855: TLS truststore initialization with FIPS-compliant crypto providers

7.3.2

Previously, in some environments using FIPS-compliant crypto providers, TLS truststore initialization failed. This occurred because of incompatible handling of empty keystore initialization. This issue is resolved by updating the truststore initialization logic to ensure compatibility with FIPS-compliant providers. This allows the TLS setup to complete successfully.

CDPD-94474: Kudu tablet server crashes on RHEL 9

7.3.2

Previously, Kudu tablet servers crashed on Red Hat Enterprise Linux 9 (RHEL 9) x86_64 platforms when the libgcc package was updated to version 11.5.0-11 or newer. These crashes typically manifested as a Segmentation Fault (SIGSEGV) or an Abort (SIGABRT) error under specific workloads. This issue is now resolved.

CDPD-80637: Kudu binary crashes on older CPU architectures

7.3.2

Previously, Kudu binaries could crash with an "Illegal instruction" error when running on older CPU architectures that did not support newer instruction sets. This occurred because the Rock sDB library, which is linked to Kudu, was built on newer hardware without the portability option enabled, leading to compatibility issues when the binaries were deployed on older machines. This issue is now resolved.

Apache Jira: [KUDU-3580](#)

CDPD-80637: NullPointerException in Kudu Java client

7.3.2

Previously, the Kudu Java client could experience a **NullPointerException** within the `TableLocationsCache.get()` method. This occurred because the client failed to verify if a cache entry specifically for Kudu master locations existed before attempting to check its status. If the entry was absent, the system would attempt to call a method on a null object, leading to a client crash.

This issue is now resolved. The `TableLocationsCache` now includes a proper null check, ensuring that the Kudu Java client handles missing cache entries gracefully without triggering an exception.

Apache Jira: [KUDU-3698](#)

CDPD-80637: Support for RSASSA-PSS signed certificates in channel bindings

7.3.2

Previously, the system could not correctly determine the hash algorithm for certificates signed with RSASSA-PSS. Because the hash algorithm in RSASSA-PSS is configurable, the existing logic failed to identify it, preventing successful channel bindings when these certificates were used.

This issue is now resolved. This ensures compatibility with modern security standards and allows for secure channel bindings when using these certificate types.

Apache Jira: [KUDU-3663](#)

CDPD-80637: Incorrect Content-Type headers for binary webserver responses

7.3.2

Previously, the Kudu webserver incorrectly identified certain binary responses, specifically the `/ipki-ca-cert-der` endpoint as `text/plain` instead of the expected `application/octet-stream`. This occurred because the webserver's internal logic relied on a boolean "styled" flag, which could only distinguish between styled and unstyled text, effectively ignoring the "binary" mode required for certificate downloads and other raw data. This issue is now resolved.

Apache Jira: [KUDU-3543](#)

CDPD-80637: Error handling for 'kudu table copy' tool

7.3.2

Previously, the kudu table copy command-line tool would crash if invalid or unexpected input was provided for the `--predicates` flag. Instead of providing a helpful error message, the tool would terminate abruptly, making it difficult for users to identify typos or syntax mistakes in their predicate definitions.

This issue is now resolved. The tool now includes improved input validation; it handles invalid predicate strings gracefully by reporting an actionable error message and exiting without a crash.

Apache Jira: [KUDU-3623](#)

Fixed Issues in Apache Knox

Review the list of Apache Knox issues that are resolved in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Knox issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

CDPD-94038: Uploading large files through HUE UI using Knox Gateway fails

Fixed an issue where uploading large files through the HUE UI using Knox Gateway failed. Large file uploads now complete successfully.

CDPD-90879: Knox service discovery does not detect Cloudera Manager cluster-level restart commands

Knox service discovery now detects cluster-level `RollingRestart` and `RestartWaitingForStalenessSuccess` commands in Cloudera Manager and automatically regenerates topologies to reflect configuration changes.

Apache JIRA: [KNOX-3194](#)

CDPD-78656: Health test for Knox fails if gateway.client.auth.needed is set to true

Previously, the health test for Knox Gateway failed if the `gateway.client.auth.needed` parameter was set to true.

This issue is now resolved. For TLS Mutual Authentication to work, you must exclude the health topology. To do this, go to [Cloudera Manager Knox Configuration](#), locate the Knox Service Advanced Configuration Snippet (Safety Valve) for `conf/gateway-site.xml` field, and add a new entry with the following parameters:

```
Name = gateway.client.auth.exclude
Value = health
```

For more information on excluding the topology, see the [Apache Knox Documentation](#).

CDPD-77233: Knox Token TTL value of -1 set to never expire

Fixed an issue where the Knox Token API raised an `UnknownTokenException` error if the lifespan value of Knox Token TTL was set to -1. Tokens with a TTL value of -1 are now correctly handled as having no expiration when verifying tokens for authentication.

Apache JIRA: [KNOX-3075](#)

CDPD-82812: HA feature not working for Rest Catalog

7.3.2.0

Previously, the Knox topology file `cdp-share-access.xml` created during Cloudera Data Sharing setup could not handle multiple Hive Metastore (HMS) nodes. In the event of a node failure, healthy nodes could not reliably take over the workload. This issue has been resolved. The Knox topology now correctly supports High Availability (HA) for the REST Catalog, ensuring proper failover between HMS nodes.

Fixed Issues in Livy

Review the list of Livy issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.0

Cloudera Runtime 7.3.2 resolves Livy issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Livy Fixed Issues](#).

CDPD-79344: Fixed Knox Gateway gzip exception with Livy HA mode

7.3.2.0

Fixed a bug in Knox Gateway's HA mode where gzip-encoded requests caused duplicated requests and unnecessary failovers. Gzip encoding is now handled more robustly, preventing the un-gzipping error and reducing unnecessary failovers.

CDPD-67187: Fixed missing JVM class imports for Spark

7.3.2.0

JVM classes are now imported during session initialization.

CDPD-63573: Zookeeper ACL support for Livy

7.3.2.0

Added Zookeeper ACL support to Livy's High Availability mode.

For more information on High Availability with Livy, refer to [Livy high availability support](#).

CDPD-49717: Fixed issue with Livy not starting if HDFS is still in safe mode

7.3.2.0

Added checks for HDFS safe mode status during Livy server restart, with retries controlled by `livy.server.hdfs.safe-mode.interval` and `livy.server.hdfs.safe-mode.max.retry.attempts`.

OPSAPS-69748: Update Ozone filesystem jar name in Livy classpath

7.3.2.0

Updated the Livy classpath to reflect the renamed Ozone filesystem jar.

Fixed Issues in Navigator Encrypt

Review the list of Navigator Encrypt issues that are resolved in Cloudera Runtime 7.3.2.0, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

The fixed issues for Navigator Encrypt in Cloudera Runtime 7.3.2.0 include all cumulative fixes from lower versions, specifically ranging from Cloudera Runtime 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes, see [Fixed Issues](#).

CDPD-89408: NavEncrypt compile issue with kernel 5.14.0-570

7.3.2

Fixed an issue where the NavEncrypt kernel module failed to compile and install on Red Hat with kernel version 5.14.0-570.35.1.el9_6.x86_64 due to incompatible inode operation definitions.

CDPD-87940: Crash in Navigator Encrypt kernel module if `--profile-file=` option is used with ACLs

7.3.2

When you added ACLs with the `acl --add --rule` command, if you included any ACLs that used the `--profile-file=` option, the kernel module (and host) crashed during Navigator Encrypt shutdown.

The issue has been fixed now.

CDPD-81638: ztrustee-c-api memory corruption getting uuids with large replies

7.3.2

Fixed an issue in the ztrustee C API where `get` `deposit` operations for UUIDs with large responses could cause memory corruption, potentially leading to runtime errors on some platforms.

Fixed Issues in Oozie

Review the list of Oozie issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Oozie issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.700. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

CDPD-93325: Hive2Main.java can print out passwords

7.3.2

Previously, Hive2Main was printing full Beeline command arguments to standard output, which might unintentionally expose LDAP credentials and the truststore password. This issue is now fixed.

CDPD-78506: Fix PATH environment variable creation for Shell action

7.3.2

With this fix, in Oozie's shell actions, the PATH environment variable is evaluated based on the YARN NodeManager host's settings. So, the PATH environment variable applies to the Launcher AM container when executed, rather than being evaluated on the Oozie server's JVM (the legacy method). You can revert to the legacy behavior by setting `oozie.action.shell.setup-path-in-oozie-server` to `true` in `oozie-site.xml`.

Also, from now, Oozie also allows defining action-specific environment variables using `oozie.launcher.<***ACTION_TYPE***>.action.env.<***VARIABLE_NAME***>` allowing you to customize execution environments for each action type in the Launcher AM.

CDPD-78069: Oozie action configuration's Java options are not applied after CDPD-60551

7.3.2

Added the ability to configure a list of Hadoop properties that can be passed to the Oozie launcher configuration, if they are not already set in Oozie through `oozie.launcher.<hadoop property key>` property or within its workflows.

The default properties in this list are `yarn.app.mapreduce.am.admin-command-opts` and `yarn.app.mapreduce.am.command-opts`. These defaults can be overridden by setting a comma-separated list in the `oozie.service.HadoopAccessorService.global.yarn.java-options-keys` property. The `oozie.LauncherConfigurationInjector.hadoop.search.properties` property is now deprecated.

CDPD-77768: Oozie log parameters are not substituted correctly

7.3.2

Previously, Oozie log parameters, for example, `{0}`, were not substituted when log messages contained apostrophes. The formatting engine treated the text following a single quote as a literal block, preventing parameter replacement. This issue is now fixed. Oozie now automatically escapes single quotes in log messages, ensuring that all parameters are substituted correctly.

CDPD-76135: Schema check database connection throws PSQLException

7.3.2

The schema check now completes successfully during Oozie startup for secure database connections or when custom connection properties are used, preventing warnings from appearing in the logs.

CDPD-68425: Oozie should validate the backend database identifiers in a case-insensitive way

7.3.2

Resolved an issue where Oozie incorrectly reported missing tables due to issues with table name casing.

Fixed issues in Ozone

Review the list of Ozone issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Ozone issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

CDPD-80567: Snapshot garbage collection issues

7.3.2

Resolved multiple issues that prevented the snapshot garbage collection system from collecting garbage. This release improves the efficiency and reliability of snapshot garbage collection, resulting in better performance and ensuring that storage is reclaimed in a timely manner.

Apache JIRA: [HDDS-12558](#)

CDPD-84361: The KeyDeletingService failed when the key size exceeded the Ratis buffer size

7.3.2

Previously, when the KeyDeletingService was fetching keys to be deleted based on keyLimitPerTask, the deletion operation failed if the key size exceeded the Ratis buffer limit (default 32 MB). This issue is fixed now and the key size does not depend on Ratis buffer size.

Apache JIRA: [HDDS-13213](#)

CDPD-80739: Ozone Recon - Containers page displayed an incorrect label for the count of blocks in various unhealthy container states

7.3.2

Previously, the Ozone Recon UI showed the Number of Keys label instead of the Number of Blocks for various unhealthy states of containers.. This issue has been fixed now and the labels are displayed correctly.

Apache JIRA: [HDDS-12588](#)

CDPD-84620: The Ozone Recon unable to serve request and throws 500 error ServiceNotReadyException on /keys/open

7.3.2

Previously, Ozone Recon would return an HTTP 500 error with a ServiceNotReadyException when the /keys/open API was called while the NSSummary tree was being rebuilt or was temporarily inconsistent. This issue is fixed now.

Apache Jira: [HDDS-13763](#)

CDPD-87883: The processed_keys_metrics table failed to updatr when converting deleted keys

7.3.2

Previously, when the Ozone tiering workflow attempted to convert the deleted keys, the conversion failed as expected; however, the failure details were not recorded in the `processed_keys_metrics` table. This occurred because deleted keys lacked required fields, such as replication type, replication factor, and so on. This issue is fixed now and the `processed_keys_metrics` table is getting updated correctly.

CDPD-69122: Ozone Manager database checkpoint generation failure

7.3.2

Previously, the Ozone Manager database checkpoint generation failed due to an `InterruptedException Unable` to process metadata snapshot request during the parallel snapshot operations and cluster restarts. This issue is fixed now.

Apache JIRA: [HDDS-10739](#)

CDPD-92017: Set ozone.om.group.rights configuration default value to ALL

7.3.2

Previously, the `ozone.om.group.rights` configuration value `READ, LIST` could not be processed by lower version of Ozone. This issue is now fixed by setting the default value of the `ozone.om.group.rights` configuration to `ALL`.

CDPD-75981: Default native ACL limits to user and user's primary group

7.3.2

Previously, the default native ACL created for an object such as volume, bucket, or file limits to the object owner and owner's primary group. If Ranger was enabled, those ACLs did not take effect, but they get saved into `KeyInfo` regardless. This issue is fixed now.

Apache JIRA: [HDDS-11656](#)

CDPD-87831: Storage Container Manager (SCM) schedules replications on full DataNodes by failing to track pending operation size

7.3.2

Previously, Storage Container Manager (SCM) scheduled replication commands to fix under-replication or misreplication for container moves, decommissioning, and other operations for both Radis and EC containers. SCM checked whether a target `DataNode` had space equal to twice the container size value before selecting it as the target node for container replication. However, SCM did not account for the pending operation size of the scheduled tasks. Consequently, SCM could over-schedule replications to a target `DataNode` that did not have enough space. This issue is now fixed.

Apache JIRA: [HDDS-13437](#)

CDPD-80178: Missing check for space availability for all DNS while container creation is in pipeline

7.3.2

Previously, if the leader node in the pipeline did not have the capacity to create a new container, it might have returned a container creation failure. If the follower node did not have the capacity to create a new container, it might have failed and repeatedly attempted to find another follower node. This behavior could cause excessive disk space consumption by parallel write blocks through a state machine, resulting in slower write performance and delayed failure responses. This issue is now fixed by checking whether a `DataNode` has enough space for a new container before allocating one. This improves write performance and reduces container creation failure in scenarios when `DataNodes` have less than 5GB disk space remaining.

Apache JIRA: [HDDS-12468](#)

CDPD-87749: No logs are available about on-demand scan triggering

7.3.2

Previously, there were no logs or debug information about why the on-demand scans were triggering on the containers. This issue is fixed now, and logs are available specifying the reason for on-demand container scans.

Apache JIRA: [HDDS-13423](#)

CDPD-85250: The OzoneTokenIdentifier does not serialize or deserialize correctly

7.3.2

Previously, the null `omServiceId` was de-serialized as an empty string, which can cause the delegation token cleanup issues in RocksDB. This issue is fixed now.

Apache JIRA: [HDDS-13264](#)

CDPD-82295: AWS S3 DeleteObject failures for FSO bucket keys containing special characters

7.3.2

Previously, AWS S3 DeleteObject could fail for FSO bucket keys containing special characters. This issue has been fixed now by removing name validation during deletion.

Apache JIRA: [HDDS-12911](#)

CDPD-74686: DirectoryDeletion task ignored through ratis

7.3.2

Previously, directory deletion task was ignored through Ratis, and there was deletion retry instead of actually deleting it. This issue is now resolved.

Apache JIRA: [HDDS-11491](#)

CDPD-74685: Directory deletion get stuck having millions of directory

7.3.2

Previously, directory deletion cleanup at background failed if there were huge number of empty directories (millions of empty directories), whose combined metadata size crossed the Ratis allowed request size for deletion. This issue is now resolved.

Apache JIRA: [HDDS-11492](#)

CDPD-87270: Secret key premature expiration and invalidation

7.3.2

Resolved the secret key issue and adjusted the secret key expiry calculation (`hdds.secret.key.expiry.duration = 9d`) to account for delegation token lifetimes. This change ensures that tokens remain valid for their full configured duration, preventing premature key invalidation and improving stable authentication.

Apache JIRA: [HDDS-13343](#)

CDPD-76523: ozone debug ldb --with-keys key defaults to false instead of true

7.3.2

Previously, the ozone debug `ldb --with-keys` option defaulted to false instead of true when specified without a value and did not print the keys. This issue is fixed now, and the ozone debug `ldb --with-keys` option now defaults to true when specified without a value so it will include keys in the output by default.

Apache JIRA: [HDDS-11782](#)

CDPD-84609: The --output-dir option was removed from replicas verify command

7.3.2

Previously, the Ozone debug `replicas verify --output-dir` was not being used for replicas verify command. In future writing the output into a specified file/ will be supported, so instead of removing the command it has to be made optional.. This issue is fixed now and the `--output-dir` option is made optional for replicas verify command and is no longer a required field.

Apache JIRA: [HDDS-13248](#)

CDPD-76520: DataNode aborts if hdds.datanode.wait.on.all.followers = true

7.3.2

Previously, the DataNode aborted if the `hdds.datanode.wait.on.all.followers` configuration was set to true. This issue is fixed and the DataNode no longer aborts if the `hdds.datanode.wait.on.all.followers` configuration is set to true.

Apache JIRA: [HDDS-11785](#)

CDPD-76501: DataNode Ratis is taking snapshots frequently

7.3.2

Previously, DataNode Ratis was taking snapshots every 5 to 8 seconds causing overhead. This issue is fixed now and the `hdds.ratis.snapshot.threshold` and `hdds.container.ratis.statemachine.max.pending.apply-transactions` configuration limits are increased to 100k to avoid frequent DataNode Ratis snapshotting.

Apache JIRA: [HDDS-11773](#)

CDPD-75112: HBase RegionServer crashes due to inconsistency caused by Ozone client failover handling

7.3.2

Previously, the HBase RegionServer crashes due to inconsistency caused by Ozone client failover handling. This issue is fixed now by making the Ozone Manager client retry idempotent so that the clients do not crash due to inconsistent results.

Apache JIRA: [HDDS-11558](#)

CDPD-77938: Local Refresh button for current selected path is missing in the new Ozone Recon UI

7.3.2

Previously, refreshing the Recon UI page would reset the current path selection and return users to the root directory, causing loss of context and requiring manual navigation. This issue is fixed now and a new Path Reload button in the new Recon UI for the Namespace page.

Apache JIRA: [HDDS-12085](#)

CDPD-77728: Calendar disappears while setting custom date range in the Heatmap page in New Recon UI

7.3.2

Previously, setting the custom date range in the Heatmap page of the new Recon UI would cause the calendar widget to close unexpectedly. Specifically, when you clicked the back arrow to navigate to a previous month in the datepicker, the entire calendar and dropdown would disappear, preventing date selection. This issue is fixed, and the calendar remains visible until a date is selected and confirmed, allowing users to set custom date ranges as intended.

Apache JIRA: [HDDS-12044](#)

CDPD-77356: Recon UI displayed identical and duplicate values for Quota Allowed and Quota In Bytes

7.3.2

Previously, in the Ozone Recon UI, the Quota Allowed and Quota In Bytes fields incorrectly displayed the same value. This duplication prevented users from accurately distinguishing between the allocated quota and the actual disk space consumed. This issue is fixed now and the Recon UI displays the values correctly.

Apache JIRA: [HDDS-11987](#)

CDPD-74437: Multiple IOzoneAuthorizer instances may be created during Ratis snapshot installation failures

7.3.2

Previously, if a failure occurred during the installation of a Ratis snapshot after the metadata manager was stopped, multiple instances of the Ozone authorizer could be created and retained in memory. This led to excessive heap usage and, in some cases, crashes due to long garbage collection pauses, especially in environments with Ranger+Ozone integration. The issue is fixed, and the old authorizer instances are properly cleaned up, preventing heap exhaustion.

Apache JIRA: [HDDS-11472](#)

CDPD-92003: Container Size Count Task showing empty in new Recon UI

7.3.2

Previously, in the Ozone Recon UI, the Container Size Count Task page appeared empty when accessed through the new user interface. This issue is fixed now.

Apache JIRA: [HDDS-13821](#)

CDPD-88628: Ozone Recon Overview page does not load until all APIs are loaded

7.3.2

Previously, the Recon Overview page waited for all API calls to complete before displaying any results, causing delays and poor responsiveness. This issue is fixed now, and each card on the Overview page now loads independently as soon as its corresponding API call resolves. This change improves overall page responsiveness and ensures that API errors only affect the relevant cards, rather than preventing the entire page from loading.

Apache JIRA: [HDDS-13542](#)

CDPD-88541: Namespace Usage page becomes blank when Recon DB is missing

7.3.2

Previously, the Namespace Usage page could appear blank if the Recon DB was missing during a fresh installation. This issue is fixed now.

Apache JIRA: [HDDS-13528](#)

CDPD-88383: Accessing the new Ozone Recon UI through Knox breaks the UI

7.3.2

Previously, accessing the new Ozone Recon UI through a reverse proxy such as Knox caused the UI to break. This issue is fixed now.

Apache JIRA: [HDDS-13512](#)

CDPD-56281: Ozone Manager database updates were blocked while Recon was reprocessing all Recon tasks

7.3.2

Previously, when Recon was reprocessing all Recon tasks, Ozone Manager database updates were blocked, which could cause repeated full snapshots and impact performance. This issue has been fixed by allowing Ozone Manager database updates to proceed concurrently with Recon task processing, preventing unnecessary full snapshots and improving system efficiency.

Apache JIRA: [HDDS-8633](#)

CDPD-77805: Improper error handling in the NSSummaryTask

7.3.2

Previously, improper error handling in the NSSummaryTask could lead to data inconsistencies in the Ozone Recon. This issue is fixed now, and ensures robust error handling in Ozone Recon.

Apache JIRA: [HDDS-12062](#)

CDPD-80826: Ozone Recon failures during the bootstrapping process

7.3.2

Previously, Ozone Recon failures during the bootstrapping process were not handle properly. This issue is fixed, and if any Ozone Manager (OM) task fails during bootstrapping, Recon will now

correctly handle and reprocess failed OM tasks to ensure successful bootstrapping. Additionally, if a partial or corrupted OM database tarball is received, Recon will clean up the corrupted file and restart the fetch process from scratch, maintaining data consistency and integrity.

Apache JIRA: [HDDS-12615](#)

CDPD-76226: The Recon ListKeys API did not return an appropriate HTTP response

7.3.2

Previously, the Recon ListKeys API did not return an appropriate HTTP response when an NSSummary rebuild was in progress. With this fix, the API now returns HTTP status code 503 (Service Unavailable) to indicate that the service is temporarily unavailable due to the ongoing NSSummary rebuild. This allows clients to properly handle the too busy or try again later scenario.

Apache JIRA: [HDDS-11708](#)

CDPD-76248: The default volume choosing policy is not updated correctly in the ozone-default.xml

7.3.2

Previously, the default volume choosing policy was updated from RoundRobinVolumeChoosingPolicy to CapacityVolumeChoosingPolicy. However the ozone-default.xml is not updated correctly and still considers the RoundRobinVolumeChoosingPolicy as the default volume choosing policy. But, the RoundRobinVolumeChoosingPolicy did not consider available space on volumes when creating or replicating containers, which could result in allocate block failures (though retried) or the creation of small containers. This issue is fixed now, and the default volume choosing policy has been changed from RoundRobinVolumeChoosingPolicy to CapacityVolumeChoosingPolicy in the ozone-default.xml. This ensures that available capacity is now taken into account during container allocation, improving reliability and resource utilization.

Apache JIRA: [HDDS-11735](#)

CDPD-73809: Multithreading issue in the ContainerBalancerTask

7.3.2

Previously, the concurrent access to shared data structures in the getCurrentIterationsStatistic method could cause unpredictable errors. This issue is fixed now, and ensures thread safety of each access inside the getCurrentIterationsStatistic method by synchronizing access to the iterationsStatistic list and using ConcurrentHashMap for concurrent access to maps from findTargetStrategy and findSourceStrategy.

Apache JIRA: [HDDS-11386](#)

CDPD-88723: The FSORepairTool fails to distinguish Unreachable and Unreferenced objects

7.3.2

Previously, the FSORepairTool logic to distinguish between Unreachable and Unreferenced objects was incorrect. This issue is fixed now, and the logic has been corrected. The unreachable objects are not marked for repair as background cleanup processes will eventually handle them, while objects that are neither reachable nor unreachable are classified as unreferenced and marked for repair.

Apache JIRA: [HDDS-13549](#)

CDPD-87575: The ozone admin container create command runs forever without kinit

7.3.2

Previously, the ozone admin container create command would run indefinitely on secure Ozone clusters with multiple Storage Container Manager (SCM) nodes if authentication failed (for example, when kinit was not performed). This issue is now fixed, and the retry logic has been updated to fail fast on authentication exceptions, providing immediate feedback to the user instead of hanging. This issue was observed only in SCM HA cluster configurations.

Apache JIRA: [HDDS-13405](#)

CDPD-90362: Container Balancer stop command fails with an error as already stopped

7.3.2

Previously, the `stopBalancer` command for the Ozone Container Balancer would fail with an error if the balancer was already stopped, instead of returning a successful response. This issue is now fixed. The `stopBalancer` operation is now idempotent and will return success if the balancer is already stopped.

Additionally, a race condition was resolved where, during an SCM leadership change, the balancer could unintentionally restart due to the persisted state not being updated. Now, the system ensures that the balancer's stopped state is correctly persisted, preventing unintended restarts during leadership transitions.

Apache JIRA: [HDDS-13694](#)

CDPD-89400: DataNode pipeline closes frequently

7.3.2

Previously, when the DataNode (DN) Ratis identified issues with a pipeline—such as a slow follower, prolonged leader election, or disk failures—it repeatedly triggered Close Pipeline actions, even if a close action was already pending in the DN command queue. This could result in excessive close actions being queued on every heartbeat, leading to inefficiency and potential command queue bloat. The issue is now fixed and introduces a check to ensure that a Close Pipeline action for a specific pipeline is not added to the command queue if one is already pending, thereby preventing redundant triggers and optimizing the signaling mechanism.

Apache JIRA: [HDDS-13618](#)

CDPD-80991: Non-administrative users could attempt to perform Ozone Manager (OM) decommission

7.3.2

Previously, non-administrative users could attempt to perform OM decommission, which could lead to unauthorized or unintended changes. This issue has now been fixed, and only users with administrative privileges are authorized to perform OM decommission actions, enhancing the security and integrity of cluster management.

Apache JIRA: [HDDS-12646](#)

Fixed Issues in Apache Parquet

Review the list of Parquet issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

CDPD-74102: Migrate Parquet to Jackson 2

7.3.2

This fix migrates Parquet from the deprecated Jackson 1 (`org.codehaus.jackson`) to Jackson 2 (`com.fasterxml.jackson`).

CDPD-74103: Migrate dependency from commons-lang (v2) to commons-lang3

7.3.2

This fix replaces all `org.apache.commons.lang.*` imports with updated `org.apache.commons-lang3.*`. This removes dependency on the deprecated commons-lang version 2 library.

CDPD-77502: Upgrade Parquet to version 1.15.1

7.3.2

The Parquet component has been upgraded from version 1.12.3 to version 1.15.1.

Fixed Issues in Phoenix

Fixed issues for Phoenix are addressed in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Phoenix issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.700. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

CDPD-76203: Fails to update to Phoenix 5.1.3 with PhoenixNonRetryableRuntimeException

7.3.2

While updating to the Phoenix 5.1.3 version, the `PhoenixNonRetryableRuntimeException` is thrown which states that the `org.apache.phoenix.query.DefaultGuidePostsCacheFactory` class cannot be loaded or instantiated from SquirrelSQL.

This issue is resolved now.

Apache Jira: [PHOENIX-7290](#)

Fixed Issues in Ranger

Review the list of Apache Ranger issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

The fixed issues for Apache Ranger in Cloudera Runtime 7.3.2 include all cumulative fixes from lower versions, specifically ranging from Cloudera Runtime 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes, see [Fixed Issues](#).

CDPD-91064: USER role accounts to access only their own user details

7.3.2

Fixed an issue where Ranger users with only the USER role could retrieve details of admin or keyadmin users with similar usernames via the `/service/xusers/users` API. The API now restricts USER-role accounts to accessing only their own user details.

CDPD-90358: Hostname verification issue

7.3.2

Fixed an issue in Ranger where hostname verification did not correctly validate the hostname across the certificate chain, improving SSL/TLS security for secured connections.

CDPD-87716: Ranger RAZ on S3 path with encoded characters

7.3.2

Fixed an issue where Ranger RAZ did not correctly evaluate S3 access policies for Iceberg table partitions when the S3 path contained encoded characters (for example, the `=` in partition directory names encoded by the AWS SDK v2). Because the encoded path was not decoded before policy evaluation, deny policies on those partition directories were not applied as expected. The path is now decoded prior to authorization checks, ensuring that deny policies on Iceberg partition directories are honored correctly.

CDPD-85478: [GDS] Dataset details not visible in Ranger access logs for RMS enforcement

7.3.2

Fixed an issue where, for RMS enforcement cases, Ranger access logs showed only the dataset policy ID but did not include the corresponding GDS dataset details. Audit entries now correctly capture the GDS dataset information for RMS-enforced datasets.

CDPD-82386: Ranger does not clean up stale plugin entries after role deletion or migration

7.3.2

Ranger now provides supported REST APIs to remove stale plugin status entries from Ranger Admin. Previously, plugin records for deleted or migrated service roles were not cleaned up and continued to appear under Audits > Plugin Status, causing confusion about plugin health and forcing risky manual database updates. The fix introduces ranger-admin endpoints to delete obsolete plugin entries either by ID or by attributes (such as service name, plugin host name, and application type), allowing administrators to safely clean up outdated plugin information and keep the Plugin Status view accurate.

CDPD-81148: Ranger enters an infinite loop and generate persistent database errors

7.3.2

Fixed an issue where Ranger could enter an infinite loop and generate persistent database errors when concurrent sessions updated or deleted the same policy, role, or tag. Tasks scheduled to run after the main transaction (for example, policy label updates, role version updates, and tag version updates) are now processed correctly, preventing infinite loops and ensuring that database errors no longer require a Ranger restart to recover.

CDPD-80625: Improved exception handling for RAZ GCP HMAC keys

7.3.2

Resolved an issue in the RAZ GCP HMAC key generation lifecycle by ignoring transient storage exceptions and retrying key retrieval a few times, and improved related logging. This makes HMAC key creation and deletion more reliable for GCP RAZ integrations.

CDPD-79459: RMS full-sync fails with unsupported Hive table location schemes

7.3.2

Previously, Ranger RMS full-sync could fail with an exception when Hive table locations were stored on file systems that RMS does not support (for example, Azure abfss), causing full-sync to stop before completion.

This issue has been fixed by skipping unsupported file schema types when processing table and database metadata during RMS full-sync and delta-sync. Ranger RMS ACL synchronization now continues for supported file schemes, and locations on unsupported file systems are not mapped; instead, an informational message is logged in the RMS server logs.

CDPD-78321: Performance fixes for Ozone plugin

7.3.2

Fixed the following performance issues observed while evaluating policies for multi-level resources:

- RANGER-4893: Improves policy evaluation for multi level resource hierarchies.
- RANGER-4922: Reduces time to find tags associated with multi-level resources.

CDPD-78151: Added configuration to control append mode for HDFS audit writes

7.3.2

A new configuration parameter was added to Ranger to control whether APPEND mode is used when writing audits to HDFS after errors or exceptions. Previously, Ranger would attempt to append to existing HDFS audit files in such cases (falling back to WRITE mode if append was not possible) to avoid generating a large number of small audit files. With this change, administrators can explicitly enable or disable the use of APPEND mode for HDFS audit writes, allowing better control over audit file handling in specific deployment scenarios.

CDPD-77949: CSV injection vulnerability during CSV/Excel export from Ranger Admin

7.3.2

Fixed an issue that could allow CSV injection during CSV and Excel exports from Ranger Admin (CVE-2024-55532). The export of Ranger policies to CSV/Excel has been removed from the Ranger Admin UI, and the affected export APIs have been deprecated.

CDPD-76633: Ranger RMS server throws ConcurrentModificationException

7.3.2

Previously, the Ranger RMS server could enter an unrecoverable error state with a ConcurrentModificationException when large service resource mappings were downloaded to the NameNode while Hive metadata changes were being applied. This could lead to follow-on errors when fetching resource#mapping deltas and no clear path to restore the RMS server to a normal state.

The fix ensures that RMS now uses a shallow copy of the service resource mappings before applying deltas, so the mappings are not modified while they are being serialized for download, preventing the ConcurrentModificationException and stabilizing RMS behavior under concurrent load.

CDPD-75532: Remove self node from the resourceTrie only if it has no children, no evaluators and no wildcard-evaluators

7.3.2

When two policies have a common subset of resources and are defined on the same user (or subset of users, through groups or direct users), if one of these policies is modified (on anything: name, resource, user), it is the only one in effect during access evaluation, until a restart of the underlying service.

This issue has been fixed now.

CDPD-74403: Fixed hardcoded parcel path and sql driver in authzmigrator

7.3.2

Fixed an issue where the authzmigrator/authz-export.sh script failed with NoClassDefFoundError : javax/jdo/JDOHelper when a custom parcel directory or non#default database driver was used. The script now uses the Cloudera Manager parcel directory configuration instead of a hard#coded Cloudera parcel path and automatically selects the appropriate JDBC driver based on the database flavor.

CDPD-73935: Fixed an issue where Ranger “federated user” accounts could log in and perform operations

7.3.2

Ranger now validates the federated user type and prevents these external, data#sharing users from logging in to Ranger, ensuring they are used only for metrics, access history, and audit purposes related to data#sharing features.

CDPD-73779: Support a new user type for external users from Data Sharing

7.3.2

Fixed an issue where Ranger did not distinguish users created by the Data Catalog external user registration process from other external users. Ranger now supports a new Federated User type in the Ranger Admin UI to represent users originating from Data Sharing (Data Catalog), in addition to the existing Internal and External user types.

CDPD-71673: Security Zone policies version increment issue

7.3.2

This fix addresses an issue where updating a resource caused the associated Security Zone policy version to increment by two instead of one.

CDPD-71563: Issue in dedupTag() method

7.3.2

Fixed a logical flaw in Ranger tag de-duplication that could incorrectly remove valid tags during the `dedupTag()` operation, preventing policy evaluation failures.

CDPD-69631: Disable Atlas service under the policy permission of tag-based policy

7.3.2

Fixed an issue where the Ranger UI incorrectly allowed selecting the Atlas service in tag-based policy permissions, even though tag-based policies are not supported for Atlas. The Atlas service option is now disabled in tag-based policy permissions to prevent misconfiguration.

CDPD-68970: Fixed an inconsistency between the Ranger UI and Policy Creation API

7.3.2

The Policy Creation API now validates input and fails policy creation if the policy contains only empty values or includes `[""]` or `["null"]` in `policyItem` users, groups, or roles, aligning API behavior with the Ranger UI.

CDPD-68500: Ranger policy create/update accepts duplicate group names

7.3.2

Previously, Ranger allowed duplicate group and role entries to be added to policies when using the public policy API (`/service/public/v2/api/policy`), even though the Ranger UI blocked such duplicates. This caused policies created or updated via the API to contain repeated group or role names. The API has been updated to remove duplicate user, group, and role entries during policy creation and update, ensuring consistent behavior with the Ranger UI.

CDPD-68297: REST endpoints do not prevent duplicate values for a resource

7.3.2

REST-based policy creation in Ranger now prevents duplicate values for a resource. Previously, REST API calls could create policies with duplicate resource values (for example, a database list like `[test_db1, test_db1]`), which could result in multiple policies for the same resource. Policy validation has been updated to reject such requests and require all values for a given resource to be unique.

CDPD-67359: Permissions issues while trying to access folders in Hue file browser

7.3.2

Fixed an issue in the Hue File Browser where accessing S3 directories with more than 1000 objects whose names contained the `"="` character failed with a `"Cannot Access: <s3 Path>"` error. The S3 marker parameter encoding has been corrected so that directories with 1000+ such keys can now be listed successfully through RAZ.

CDPD-67269: Support multiple resource sets in a policy

7.3.2

Improved Ranger policy evaluation to fully support multiple resource sets within a single policy, aligning Cloudera Ranger behavior with the upstream Apache Ranger RANGER-3796 enhancement.

CDPD-62008: Ranger ABAC now supports internal user and group attributes

7.3.2

Resolved an issue where Ranger ABAC policies could not leverage certain internal user and group attributes. ABAC policies can now use the following internal attributes for access control, masking, and row-filtering decisions: `syncSource`, `isInternal`, and `emailAddress`.

CDPD-40734: User allowed to insert data into a hive table when there is a deny policy on a table column

7.3.2

A user is allowed to enter data into a table even if there is a deny policy present on one of the table columns.

This issue has been fixed now.

OPSAPS-75602: Issue with RANGER_C719 CSD becoming stale after upgrading Cloudera Manager

7.3.2

Fixed an issue where the RANGER_C719 CSD could become stale after upgrading Cloudera Manager from 7.13.1.600 with Cloudera 7.1.9 to 7.13.2.0 by fixing the following:

- OPSAPS-73498: Added Cloudera Manager side ranger-trino integration changes.
- OPSAPS-73152: Improved Ranger Admin Diagnostic collection command from Cloudera Manager scripts.

OPSAPS-75556: After upgrade from 7.1.9 to 7.3.2.0 dataset field type is set to boolean in solr managed-schema

7.3.2

Fixed an issue where, after upgrading from Cloudera 7.1.9 to 7.3.2, the datasets field in the ranger_audits Solr collection schema was incorrectly set to the boolean type instead of key_lower_case with multiValued="true". This schema mismatch caused Ranger Admin to fail to load the Access Audit page on upgraded clusters. The upgrade process now updates the ranger_audits Solr schema so that the datasets field is created with the correct type and behaves consistently with fresh 7.3.2 deployments.

OPSAPS-71619: Removed the mandatory validation for ranger.ldap.user.dnpattern

7.3.2

Previously, when LDAP was configured as the external authentication type for Ranger Admin, the ranger.ldap.user.dnpattern parameter was mandatory. If it was not set, the Ranger Admin service failed to start, even though this parameter is rarely required and is ignored when LDAP bind DN/password and user search parameters are configured. This has been fixed by removing the mandatory validation for ranger.ldap.user.dnpattern, so the parameter is now optional and the service can start without requiring a dummy value.

OPSAPS-69156: Fixed an issue with Java add-opens/add-modules/add-exports options

7.3.2

Cloudera Manager components now consistently use the --add-opens=, --add-modules=, and --add-exports= syntax for Java options. This avoids cases where options passed via JAVA_TOOL_OPTIONS could be rejected (for example when using --add-opens or --add-exports without =), improving compatibility across different Java runtimes.

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

7.3.2

Previously, Cloudera Manager reported the Ranger RMS server as healthy based only on the RMS process (PID), even when the RMS web service was not fully initialized and the service was inaccessible. The health check logic has been updated to use a Cloudera Manager web alert that verifies the Ranger RMS web endpoint instead of relying solely on the PID. This allows Cloudera Manager to more accurately detect when RMS is not accessible and helps users identify RMS availability issues faster.

Fixed Issues in Ranger KMS

Review the list of Ranger KMS issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

The fixed issues for Ranger KMS in Cloudera Runtime 7.3.2 include all cumulative fixes from lower versions, specifically ranging from Cloudera Runtime 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes, see [Fixed Issues](#).

CDPD-74162: Missing HSTS Headers for 404 Not found requests port 9494

7.3.2

Fixed an issue where Ranger KMS HTTPS endpoints on port 9494 did not return HTTP Strict-Transport-Security (HSTS) headers for 404 Not Found responses. Ranger KMS now consistently includes HSTS headers on 404 responses, improving security for clients accessing the KMS service.

CDPD-80814: Failed to migrate the Ranger Kms master key if the masterkey is stored with HDP format

7.3.2

Fixed an issue where migrating the Ranger KMS master key from an HDP-formatted master key to HSM failed with a `NullPointerException`, preventing import of the master key from the Ranger database. The migration utility now correctly handles the older HDP master key format and completes the migration successfully.

OPSAPS-72766: Ranger KMS tomcat context update

7.3.2

Updated the default Tomcat context for Ranger KMS from `/kms` to `/` by changing the `ranger.contextName` property in `ranger-kms-site.xml`. This aligns the Ranger KMS context path with Cloudera configuration and simplifies access and integration.

Fixed Issues in Schema Registry

Fixed issues and resolved maintenance items for Schema Registry are addressed in Cloudera Runtime 7.3.2 and its associated service packs.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Schema Registry issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

OPSAPS-71866: Schema Registry fails to start when `schema.registry.oauth.clock.skew` is set in Cloudera Manager

7.3.2

Schema Registry failed to start when the `schema.registry.oauth.clock.skew` configuration property was set in Cloudera Manager. This issue is now fixed.

OPSAPS-71679: Ranger resource-based policy creation attempted on each service restart

7.3.2

On each service restart, Kafka, Kafka Connect, and Schema Registry attempted to create Ranger resource-based policies even when the Ranger policy cache was already present. Resource-based policy creation is now skipped when the policy cache file exists, reducing unnecessary Ranger operations on restart.

CDPD-73277: Schema Registry configuration with Oracle TLS fails

7.3.2

The TLS-related Oracle database configuration parameters no longer cause issues during the Schema Registry startup if they are set.

CDPD-94371: Schema Registry does not start with MySQL 8.4 on new installations

7.3.2

Fixed an issue that caused startup failure on new Schema Registry installations when using MySQL 8.4.

Fixed Issues in Apache Solr

Review the list of Solr issues that are resolved in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Solr issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

There are no fixed issues in this release.

Fixed Issues in Spark

Review the list of Spark issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Spark issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Spark Fixed Issues](#).

CDPD-90482: Jersey Library update

7.3.2.0

Updated the Jersey suite of libraries to versions 2.0 and above to ensure JDK 17 compatibility.

CDPD-82979: Added Cloudera-specific configuration for MapPartitionsRDD default determinacy

7.3.2.0

To avoid failures during error handling, a Cloudera-specific configuration was added to change the default determinacy level for any MapPartitionsRDD.

CDPD-80239: Fixed handling of non-deterministic SQL expressions

7.3.2.0

Fixed a bug related to intermittent data corruption and record loss during task failures. Spark now honors the `deterministic=false` attribute for non-deterministic data types.

CDPD-67187: Fixed missing JVM class imports for Spark

7.3.2.0

JVM classes are now imported during session initialization.

OPSAPS-74346, OPSAPS-74346, OPSAPS-74316: Enhancing Spark security

7.3.2.0

Changed the generation of the `spark.yarn.historyServer.address` value to use the HTTPS address when SSL/TLS is enabled. The `spark3.network.crypto.enabled` new configuration property is now available to enable AES-based encryption.

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

7.3.2.0

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

For more information, see *New Cloudera Manager configuration parameter `spark_pyspark_executable_path` has been added to the Livy for Spark 3* in [Behavioral Changes In Cloudera Manager 7.13.2](#).

OPSAPS-76539: Cloudera Manager UI allowed adding multiple Spark 3 service instances within a single cluster

7.3.2.0

Prevented multiple Spark 3 service instances within a single cluster by implementing a `maxInstances` flag in Cloudera Manager.

OPSAPS-74281: Disabled the live Spark UI when the `ENCRYPT_ALL_PORTS` feature flag is enabled

7.3.2.0

Enhanced the security of Spark by implementing new default configuration settings:

1. `spark.ui.enabled=false`: preventing initiation of an HTTP service that can be accessed from external hosts.
2. `spark.io.encryption.keySizeBits=256`: the default Spark `keySizeBits` has been increased from 128 to 256

Fixed Issues in Spark Atlas Connector

Review the list of Spark Atlas Connector issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Spark Atlas Connector issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Spark Atlas Connector Fixed Issues](#).

CDPD-95934: Added a new configuration parameter in Spark Atlas Connector to enable or disable column lineage processing.

7.3.2.0

The configuration parameter `ATLAS_SPARK_COLUMN_LINEAGE_ENABLED` is available to enable or disable column lineage processing.

Fixed Issues in Apache Sqoop

Review the list of Sqoop issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

There are no new fixes introduced in this release.

Fixed Issues in Streams Messaging Manager

Fixed issues and resolved maintenance items for Streams Messaging Manager are addressed in Cloudera Runtime 7.3.2 and its associated service packs.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Streams Messaging Manager issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

OPSAPS-75519: Streams Messaging Manager Rest Admin Server fails to start when the keystore or truststore password contains special characters

7.3.2

The Streams Messaging Manager Rest Admin Server did not handle special characters in keystore or truststore passwords correctly in the generated configuration, which prevented the role from starting. This issue is fixed.

OPSAPS-72814: Streams Messaging Manager used incorrect Ranger policy cache directory

7.3.2

Streams Messaging Manager used an incorrect Ranger policy cache directory setting, which led to permission errors when the Kafka Ranger plugin wrote its policy cache. This issue is fixed.

CDPD-70906: NullPointerException no longer displayed when user is not authorized to access Cruise Control

7.3.2

A message with the text `NullPointerException` is no longer displayed if the logged-in user is not authorized to access Cruise Control.

CDPD-74620: Data Explorer returns 500 Server Error when switching partitions

7.3.2

In the Streams Messaging Manager Data Explorer, switching between partitions now waits for the new partition's offset information before fetching data. This eliminates the 500 Server Error that occurred when the previous partition's offset was incorrectly used.

CDPD-76375: Topic column name is incorrect in the replications table

7.3.2

Previously, on the **Cluster Replications** page, in the replications table, the topic column name was **Source Topic Name** instead of **Target Topic Name**.

This issue is now fixed and the table name is the correct **Target Topic Name**.

CDPD-91339: Connector tasks that belong to the same worker cannot be selected individually

7.3.2

Fixed an issue where checking a checkbox in the **Tasks** section on a **Connector Profile** page sometimes selected multiple tasks. Now selecting a single task always selects a single task.

CDPD-62315: Add error boundary to catch runtime errors

7.3.2

An error page fallback is added for unexpected runtime UI errors.

CDPD-45542: Excessive decimal places displayed in metrics

7.3.2

Fixed an issue where task metric percentages in the **Tasks** section of a **Connector Profile** page displayed excessive decimal places, causing UI overflow. Metrics now display with one decimal place.

CDPD-88488: The Streams Messaging Manager server fails to start if its keystore password starts with a special character or includes quotation marks

7.3.2

Fixed an issue where using special characters in keystore passwords made the Streams Messaging Manager fail to start.

Fixed Issues in Streams Replication Manager

Fixed issues and resolved maintenance items for Streams Replication Manager are addressed in Cloudera Runtime 7.3.2 and its associated service packs.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Streams Replication Manager issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

OPSAPS-72697: Invalid Streams Replication Manager configuration when Cloudera Manager retried start after a failed attempt

7.3.2

Fixed an issue that could result in configuration errors (invalid configuration files) for Streams Replication Manager if Cloudera Manager automatically tried to start the service again when the first try was unsuccessful. This issue is fixed.

Fixed Issues in YARN and YARN Queue Manager

Fixed issues and resolved maintenance items for YARN and YARN Queue Manager are addressed in Cloudera Runtime 7.3.2 and its associated service packs.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves YARN and YARN Queue Manager issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.700. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

CDPD-49702: Error in NodeManager when executing /var/lib/yarn-ce/bin/container-executor

7.3.2

Previously, a job failure occurred because NodeManager returned a No such file or directory error when attempting to run the /var/lib/yarn-ce/bin/container-executor program. This issue is now resolved and NodeManager is now marked as unhealthy and shut down if it is unable to run the program.

Apache Jira: [YARN-11709](#)

COMPX-13401: CapacityScheduler UI queue filter does not work as expected when submitting applications with leaf queue's name

7.3.2

Previously, when submitting applications to YARN using only a leaf queue name, for example, default or custom, instead of the full queue path, for example, root.default, the ResourceManager (RM) and CapacityScheduler UI inconsistently displayed or filtered applications. This led to confusion, as the same queue could be displayed under different names, and applications were not visible under the expected queue filter in the UI.

This issue is now resolved and RM now returns the full queue path regardless of whether the application was submitted with a leaf queue name or a full queue path.

Apache Jira: [YARN-11538](#)

COMPX-14637: Missing permissions on NodeManagers local directories on startup

7.3.2

Previously, the NodeManager created its required local directories on startup with the correct 755 permissions only if they did not already exist.

If an administrator created these directories with incorrect permissions, or if the permissions were altered after the NodeManager started, the NodeManager failed to reset them. This lack of permission enforcement caused container failures. This issue is now resolved.

Apache Jira: [YARN-11703](#)

COMPX-17261: NodeManager disk health status oscillation with MB-based limits

7.3.2

Previously, the NodeManager disk health status could oscillate rapidly between good and full states when using the `yarn.nodemanager.disk-health-checker.min-free-space-per-disk-mb` property because it relied on a single threshold. This issue is now fixed. To prevent this oscillation, the `yarn.nodemanager.disk-health-checker.min-free-space-per-disk-watermark-high-mb` optional configuration property is now available. This setting allows administrators to specify the minimum free space in megabytes required for a previously bad disk to be marked as good again. If this value is not set or is lower than the `YARN.nodemanager.disk-health-checker.min-free-space-per-disk-mb` value, it defaults to the same value as the existing minimum free space configuration. This update applies to both `YARN.nodemanager.local-dirs` and `YARN.nodemanager.log-dirs` directories, providing more granular control over disk health checks.

Apache Jira: [YARN-9914](#)

COMPX-18004: Metrics missing in the RM UI2

7.3.2

The Replication Manager (RM) UI2 did not display metrics with the same granularity as RM UI (UI1). This made analyzing and debugging the scheduler behavior difficult, often requiring the retrieval of information from the UI. This issue is now resolved and the missing metrics are now available in RM UI2.

COMPX-18545: Setting maximum-application-lifetime using AQCv2 templates does not apply on the first submitted application

7.3.2

Setting the `maximum-application-lifetime` property using the AQC v2 templates did not apply to the first submitted application but was applied to the subsequent ones. This issue is now resolved.

Apache Jira: [YARN-11708](#)

COMPX-18909: NodeManager marked as unhealthy if an application is terminated

7.3.2

By design, Node Managers are marked unhealthy if an unrecoverable configuration error occurs, that is, the container-executor script is missing. Previously, a false positive marking occurred if an application application was terminated just before one of its containers was tried to access the localizer syslog file. This caused an `IOException`, and the Node Managers was incorrectly marked unhealthy. This issue is now resolved. The error checking is more specific, preventing these false positive unhealthy markings.

Apache Jira: [YARN-11753](#)

COMPX-20253: Missing Secure Sockets Layer (SSL) cipher inclusion list support

7.3.2

Previously, support for configuring an inclusion list of SSL ciphers for `HttpServer2` and `SSLFactory` was missing. This issue is now resolved. When an inclusion list is set, only the listed ciphers are allowed, and any cipher present in both the inclusion and exclusion lists are excluded.

Apache Jira: [HADOOP-19546](#)

COMPX-21537: Upgraded the Jersey version

7.3.2

Upgraded the Jersey framework from version 1.19 to 2.46 to enhance security, improve performance, and ensure compatibility with modern standards and dependencies.

COMPX-23154: Unnecessary direct dependency on Jersey 1 JSR311 APIs fro HTTP headers

7.3.2

Previously, direct dependency on Jersey 1 `jsr311` API for HTTP header handling relied on the `javax.ws.rs.core.HttpHeaders` class. This issue is now resolved. All usages of `javax.ws.rs.core.Htt`

pHeaders are replaced with internal constants or alternative classes, successfully removing the direct dependency.

Apache Jira: [HADOOP-19077](#)

COMPX-23191: Null Pointer Exception in Delegation Token Renewer causes all subsequent applications to fail

7.3.2

Previously, any uncaught exception in `DelegationTokenRenewer.RenewalTimerTask#run` caused all subsequent YARN applications to fail with a `java.lang.IllegalStateException: Timer already cancelled` exception. This issue is now resolved, and such failures are prevented.

COMPX-24259: Incorrect permissions on NodeManager local directories cause container failures

7.3.2

The NodeManager creates the configured local directories with 755 permissions on startup if the directories did not exist. However, if these permissions were changed after startup or if an administrator created the directories with incorrect permissions before starting YARN, the NodeManager did not reset the permissions, resulting in container failures. This issue is now resolved.

COMPX-21461: The `queuemanager_includedCipherSuites` property fails when using comma as a separator

7.3.2

Previously, the `queuemanager_includedCipherSuites` property in Queue Manager only supported colon (:) as a separator for cipher suites. When comma (,) was introduced as an additional separator, configurations using commas caused failures in property parsing.

This issue is now resolved by updating the property parsing logic to accept both colon and comma as valid separators.

COMPX-22209: Missing centralised Apache HttpComponents libraries

7.3.2

Previously, the cpx component did not use the centralized Apache HttpComponents libraries (`httpcore`, `httpclient`, `httpcore5`, `httpclient5`, `httpcore5-h2`). This issue is now resolved. The component now uses these centralized libraries, to align with internal standards and incorporate the latest fixes.

COMPX-22213: Missing centralised Bouncy Castle (`org.bouncycastle`) libraries

7.3.2

Previously, the cpx component did not use the centralized Bouncy Castle `org.bouncycastle` library versions defined in CDPD (`bcprov-jdk18on`, `bcpkix-jdk18on`, and `bcutil-jdk18on` updated from 1.78 to 1.78.1). This issue is now resolved. The component now uses these centralized libraries to align with internal dependency standards and incorporate the latest security and bug fixes.

COMPX-23423, COMPX-22836: Apache Commons Lang upgraded to 3.18.0

7.3.2

The Apache Commons Lang package is now upgraded to version 3.18.0 in Queue Manager.

Fixed CVE-2025-28924

7.3.2

Few Hadoop API endpoints are now removed from Cloudera Runtime 7.3.2. This change is a result of the fix for [CVE-2025-48924](#) and YARN's migration to common-lang 3. For more information see, <https://docs.cloudera.com/cdp-private-cloud-base/7.3.2/private-release-notes/topics/rt-pvc-api-compat-changes-hadoop.html>

Apache Jira: [YARN-10772](#)

Fixed Issues in ZooKeeper

Review the list of ZooKeeper issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves ZooKeeper issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.700. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see [Fixed Issues](#).

CDPD-92113: Zookeeper client uses IPv6 address in case of dual#stack cluster

7.3.2

Previously, in some dual-stack environments, the Zookeeper client failed to connect if IPv6 was not reachable. This issue is now resolved to improve the connection reliability for affected configurations.

CDPD-93023, CDPD-91630: ZooKeeper build broken in JDK17 line

7.3.2

Previously, the ZooKeeper JDK17 build failed due to the missing jline dependency. This issue is now resolved.

CDPD-68499: Upgraded jline version to 3.25.1

ZooKeeper is now updated to use the jline version 3.25.1, which addresses CVE-2023-50572.

OPSAPS-76345: Cloudera Manager could not connect to the ZooKeeper JMX port when JMX TLS was enabled

7.3.2

Previously, an issue triggered by JDK upgrades related to CVE-2026-21925, prevented Cloudera Manager from securely connecting to the ZooKeeper JMX port when JMX TLS was enabled (the default for secure clusters). The problem was caused by a change in the JVM. This issue is now resolved.



Note: Upgrading to a Cloudera Manager version with this fix might lead to temporary staleness.

OPSAPS-75121: TLS Protocol and Cipher Selection support for ZooKeeper

7.3.2

Previously, when two options were removed from Cloudera Manager, related to ZooKeeper TLS settings, testing showed that the TLS version and cipher list settings became inconsistent and insecure. This issue is now resolved and the Zookeeper TLS version and cipher list settings are now removed.

OPSAPS-73956: Support IPv6 access for client

7.3.2

The ZooKeeper startup script, zkserver.sh now uses the configurable variable IP_VERSION instead of a hardcoded value. This change allows the script to dynamically set the appropriate network stack preference, enabling support for IPv4, IPv6, or dual-stack environments as configured. No further manual flag adjustments are needed for IPv6 or dual-stack modes; the system automatically handles the correct settings.

OPSAPS-73624: Data Lake upgrade failure due to ZooKeeper StartPreservingDatastore error

7.3.2

Previously, during the Data Lake (DL) Zero Downtime Upgrade (ZDU) process, the upgrade failed at the ZooKeeper service upgrade phase. This blocked the upgrade and required manual intervention to proceed.

This issue in the ZooKeeper start command is now resolved. The upgrade process now handles ZooKeeper server startups more reliably, preventing the configuration mismatch error and allowing the upgrade to complete without manual steps.

OPSAPS-72239: Failure when installing Cloudera 7.3.1 using Cloudera Manager 7.13.1.0

7.3.2

Previously, when selecting the ZooKeeper gateway role during the initial installation of a Cloudera 7.3.1 cluster using Cloudera Manager 7.13.1.0 caused the wizard to fail with an internal error, preventing the cluster setup.

This issue is now resolved and Cloudera Manager now filters the ZooKeeper gateway role during the initial installation wizard, preventing the error and allowing normal cluster completion.

OPSAPS-63656: Duplicate chart displayed in the ZooKeeper UI

7.3.2

Previously, the Outstanding Requests Across Servers chart was displayed twice in the Cloudera Manager ZooKeeper Server Aggregates UI.

This issue is now resolved and the duplicate ZooKeeper chart is now removed.

OPSAPS-57641: ZooKeeper Authentication Enforcement renders ZooKeeper Canary unable to connect

7.3.2

When ZooKeeper authentication enforcement was enabled, the Cloudera Manager Service Monitor's ZooKeeper Canary could not authenticate to ZooKeeper. This caused the canary checks to fail and blocked successful installation and monitoring in secured environments.

This issue is now resolved and Kerberos-based authentication is now enabled for the ZooKeeper Canary in the Service Monitor. The canary can now connect and run health checks correctly when ZooKeeper authentication is enforced.

OPSAPS-57019: No option to enforce client authentication

7.3.2

Previously, ZooKeeper did not provide an option in Cloudera Manager to enforce client authentication, limiting the ability to meet stricter security and compliance requirements.

This issue is now resolved and a ZooKeeper configuration option added in Cloudera Manager allows administrators to require client authentication.

OPSAPS-76344: ZooKeeper service canary failed with a certificate exception

7.3.2

Previously, in Cloudera Manager 7.13.2.0, ZooKeeper service canary failed due to a `java.security.cert.CertificateException` error. This issue is now resolved.

OPSAPS-75265, OPSAPS-75418: Mismatch in ZooKeeper component version

Previously, the ZooKeeper component version displayed and used by Cloudera Manager was not aligned with the updated ZooKeeper rebase. This issue is now fixed and Cloudera Manager now uses the correct ZooKeeper version.

Apache patch information:

- [ZOOKEEPER-2590](#)
- [ZOOKEEPER-2623](#)
- [ZOOKEEPER-3100](#)
- [ZOOKEEPER-4026](#)

- [ZOOKEEPER-4236](#)
- [ZOOKEEPER-4293](#)
- [ZOOKEEPER-4293](#)
- [ZOOKEEPER-4393](#)
- [ZOOKEEPER-4604](#)
- [ZOOKEEPER-4604](#)
- [ZOOKEEPER-4728](#)
- [ZOOKEEPER-4758](#)
- [ZOOKEEPER-4762](#)
- [ZOOKEEPER-4787](#)
- [ZOOKEEPER-4787](#)
- [ZOOKEEPER-4787](#)
- [ZOOKEEPER-4843](#)
- [ZOOKEEPER-4846](#)
- [ZOOKEEPER-4876](#)
- [ZOOKEEPER-4886](#)
- [ZOOKEEPER-4889](#)
- [ZOOKEEPER-4900](#)
- [ZOOKEEPER-4909](#)
- [ZOOKEEPER-4919](#)
- [ZOOKEEPER-4933](#)
- [ZOOKEEPER-3731](#)
- [ZOOKEEPER-3860](#)
- [ZOOKEEPER-4240](#)
- [ZOOKEEPER-4415](#)
- [ZOOKEEPER-4546](#)
- [ZOOKEEPER-4753](#)
- [ZOOKEEPER-4764](#)
- [ZOOKEEPER-4799](#)
- [ZOOKEEPER-4850](#)
- [ZOOKEEPER-4860](#)
- [ZOOKEEPER-4906](#)

Known Issues In Cloudera Runtime 7.3.2

This topic describes known issues and workarounds in this release of Cloudera Runtime.

Known Issues in Apache Atlas

Known issues and technical limitations for Atlas are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2

OPSAPS-76330: Missing request/process context ID in Atlas logs after migration from log4j2 to logback using CM default pattern

7.3.2 and its SPs and CHFs

After upgrading Apache Atlas logging from log4j2 to logback, and using the default logback pattern provided by Cloudera Manager, Atlas logs no longer consistently include the request/process context ID (for example, `etp<timestamp>-<pid> - <uuid>`). This results in a regression compared

to earlier behavior with log4j2, where the context ID was consistently present for request-scoped operations. The missing context information makes troubleshooting and tracing individual requests more difficult.

None

CDPD-91624: Performance degradation: Atlas Replication is 3-4x slower in 7.3.2 compared to 7.1.9

7.3.2

A performance degradation has been observed in the Atlas table replication process after upgrading from version 7.1.9 to 7.3.2. The average time to replicate a single table has increased by approximately 300-400%. In version 7.3.2, the replication takes 3-4 seconds per table, compared to approximately 0.94 seconds per table in version 7.1.9 (calculated from 0.11s export + 0.83s import). This significant slowdown affects the time required to complete Atlas replication operations for large-scale table replication, with tests showing that only 24,000 out of 100,000 tables were replicated in more than 40 hours.

None

OPSAPS-75853: Composite replication (Atlas + Iceberg) history entries recorded with "Partial success"

7.3.2 and its SPs and CHFs

When running composite replication operations that include both Atlas metadata and Iceberg data replication, the replication history entries are recorded with a status of Partial success even when the replication completes successfully. This can cause confusion about the actual status of replication operations.

Verify the actual replication status by checking the detailed logs and target system to confirm whether the replication completed successfully despite the Partial success status.

None

CDPD-96774: Entities not created in Atlas when Hive table created on top of OFS bucket/volume

7.3.2 and its SPs and CHFs

When a Hive table is created with data stored directly on an Ozone File System (OFS) bucket or volume (without a key path), the corresponding entities are not created in Apache Atlas. Only Hive tables created with OFS key-based paths (for example, `ofs://serviceId/volume/bucket/key`) are properly tracked in Atlas. Tables pointing directly to buckets (for example, `ofs://serviceId/volume/bucket`) or volumes (for example, `ofs://serviceId/volume`) do not generate Atlas entities, preventing metadata tracking and governance for these storage configurations.

Apache JIRA: [ATLAS-5034](#)

None

CDPD-95669: React UI: Entity type clicked in Basic Search results redirects to Home page instead of triggering type search

7.3.2 and its SPs and CHFs

In the React-based Atlas user interface, when you perform a basic search that returns multiple entities and click an entity type in the Type column of the search results table, the interface incorrectly redirects to the Home page instead of remaining on the Basic Search page and triggering a new search filtered by that entity type. This navigation behavior breaks the expected workflow and makes it difficult to refine searches based on entity types directly from the search results.

Use the Advanced Search feature to filter results by entity type, or manually enter the type filter in the Basic Search interface before performing the search.

Apache JIRA: [ATLAS-5172](#)

CDPD-95658: React UI: isIncomplete Property not visible in React UI for shell entities

7.3.2 and its SPs and CHFs

When viewing shell entities in the React user interface, the `isIncomplete` property is not displayed, even though the same entity shows this property correctly in the Classic user interface. This inconsistency prevents users from identifying incomplete entities when using the React interface. Shell entities are placeholder entities created with minimal metadata, and the `isIncomplete` property indicates that these entities require additional information to be complete. The React user interface now displays the `isIncomplete` property consistently with the Classic user interface, allowing users to identify shell entities regardless of which interface they use.

None

Apache JIRA: [ATLAS-5170](#)

CDPD-63922: RowTooBigException when reading and writing data to HBase

7.3.2 and its SPs and CHFs

When Atlas performs operations that read or write large amounts of data to HBase through JanusGraph, a `RowTooBigException` can occur if the row size exceeds the configured maximum row size limit (`hbase.table.max.rowsize`). This exception has been observed in multiple escalations during operations such as patch application, advanced search, and entity creation. The error indicates that the row data being processed is larger than the allowed maximum (typically 1GB or 2GB depending on configuration). While increasing the `hbase.table.max.rowsize` property can help in some cases, it may also cause out-of-memory errors on the HBase side. The root cause is related to how JanusGraph stores data in HBase and requires deeper investigation of Atlas data storage patterns.

As a temporary workaround, you can increase the `hbase.table.max.rowsize` property value. However, be aware that significantly increasing this value may lead to out-of-memory errors in HBase. Monitor HBase memory usage after making this change.

CDPD-94690: React UI: Classification searches parent classification instead of child

7.3.2 and its SPs and CHFs

When using the React user interface to search for entities by classification, clicking the Search option on a sub-classification incorrectly triggers a search for the parent classification instead of the selected sub-classification. This causes users to see search results for the wrong classification hierarchy level, making it difficult to find entities tagged with specific sub-classifications. The search now correctly uses the selected sub-classification, returning accurate results that match the user's intended search scope.

None

Apache JIRA: [ATLAS-5168](#)

CDPD-92014: React UI: Fix UI layout issue for tables when no records are present

7.3.2 and its SPs and CHFs

When viewing tables in the React UI Search and Business Metadata sections with no records, the table headers appear visually constrained or misaligned. This layout inconsistency affects the user interface appearance and makes empty result screens look unprofessional. The table headers now display with proper alignment and layout even when no records are present, ensuring a consistent and polished user experience across all data states.

None

Apache JIRA: [ATLAS-5151](#)

CDPD-91961: Import failure due to missing support for handling multiple relationship types under a single attribute

7.3.2 and its SPs and CHFs

When importing entities where a single relationship attribute (such as `hive_db.tables`) contains a mix of different entity types (for example, both `hive_table` and `iceberg_table` references), the import operation may fail or produce incomplete relationship graphs. Atlas incorrectly derives only

one relationship type for the entire collection, typically `hive_table_db`, which causes relationship validation to fail for other entity types like `iceberg_table`. This data-dependent issue is difficult to diagnose and can result in import errors or missing Iceberg table relationships. Atlas logs may show errors about invalid relationship end types or invalid relationship definitions for specific elements within a relationship attribute.

None

Apache JIRA: [ATLAS-5156](#)

CDPD-87879: Moving an Iceberg table from one database to another database is creating incomplete entity

7.3.2 and its SPs and CHFs

When renaming an Iceberg table to move it from one database to another using the ALTER TABLE statement with RENAME TO clause (for example, ALTER TABLE test_share.iceberg_table RENAME TO default.iceberg_table), an incomplete entity is created in the target database. This occurs because the Hive Metastore hook incorrectly classifies the operation as an ADDPROPS (add properties) operation instead of a RENAME operation for Iceberg tables. This misclassification causes the HMS hook to process the event incorrectly, passing only the DO_NOT_UPDATE_STATS = true property while omitting critical properties such as old_table_name, alterTableOpType = RENAME, and old_db_name that are required for proper entity handling. As a result, Atlas creates an incomplete entity in the target database rather than properly renaming the existing entity.

None

CDPD-99217: Incomplete Metadata Display for Ozone Buckets in Atlas Search Results

7.3.2 and its SPs and CHFs

While searching for certain Ozone buckets in the Atlas UI, the results page loads, but key details such as Properties, Classifications, Relationships, and Audits are not displayed. The issue stems from the large size of the response payload, specifically the relationships section, which exceeds 500MB in the JSON response. This overwhelms the UI, preventing it from rendering associated metadata.

None

Apache JIRA: [ATLAS-5184](#)

CDPD-91959: Entities not created in Atlas when hive table created on top of a OFS bucket/volume

7.3.2 and its SPs and CHFs

In Atlas, key based tables are created properly. However, bucket/volume based tables do not show up in Atlas. When trying to create hive tables on top of ozone OFS buckets/volumes, the entities are not created.

None

Apache JIRA: [ATLAS-5034](#)

CDPD-89476: need to enhance the repair_index.py script

7.3.2 and its SPs and CHFs

When executing the repair_index.py script provided by Cloudera to reindex Atlas metadata, the script fails to run under Java 17, which is now commonly used in customer environments. Customers upgrading their OS or JVM to OpenJDK 17 encounter errors related to java.lang.reflect.InaccessibleObjectException. Even after bypassing this with additional JVM flags, the script later fails with Keystore was tampered with, or password was incorrect. This indicates potential mismatches with the Java version, keytool compatibility, or the way the script selects runtime process directories and handles credentials.

Short-term workaround:

- Add the following JVM options in `repair_index.py` to bypass JPMS restrictions in Java 17:

```
--add-opens=java.base/java.lang=ALL-UNNAMED \
--add-opens=java.base/java.lang.reflect=ALL-UNNAMED
```

- Manually verify and specify the correct conf and keystore directory to avoid loading stale credentials.
- Run the script under Java 8 or Java 11 if possible.

CDPD-88306: Atlas UI: Iceberg tables are not visible in the Relationships tab

7.3.2 and its SPs and CHFs

On the Atlas user interface, the **Relationship** tab for a database entity is not displaying Iceberg tables when both Hive and Iceberg tables are present in the same database. While Hive tables are visible as expected, Iceberg tables are missing from the visualization. If only Iceberg tables exist in the database, they are displayed correctly in the **Relationships** tab.

None

CDPD-78683: Standalone Atlas replication fails when unexpected inputs are provided

7.3.2 and its SPs and CHFs

When a standalone Atlas replication policy is created for NULL or `hive_db` entity types, and no Hive or Iceberg tables are created, the replication fails with a generic exception (Atlas Export operation has failed. Please check role logs for more info.) instead of either skipping the operation (NOOP) or providing a descriptive error message about the problematic inputs. This generic error message can be misleading.

None

CDPD-76686: Concurrent Ingestion: doesn't commit the kafka offset

7.3.2 and its SPs and CHFs

When concurrent processing is enabled (`atlas.notifications.concurrent=true`), Atlas processes messages efficiently but may fail to commit the Kafka offset. This results in the current offset getting stuck for an extended period, despite the messages being processed successfully.

None

CDPD-77640: Inconsistent results from global search due to case sensitivity

7.3.2 and its SPs and CHFs

Global search in Atlas produces inconsistent results based on the case of the search term. Searching with a lowercase term returns the expected entities, whereas searching with the same term beginning with an uppercase character does not return the same results. This issue is related to case-sensitivity handling in the search layer, likely in indexing or query processing.

None

Known issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

7.3.2 and its SPs and CHFs, 7.3.1.700 through 7.3.1.100, 7.1.9 SP1 CHF15 through 7.1.9 SP1 CHF4

Atlas hook authorization error causes `HiveCreateSysDb` command to time out due to repeated retries.

None

CDPD-69150: Unable to add labels or user defined properties in Japanese

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Adding Japanese labels or user defined properties results in the error message: “Invalid label: ###, label should contain alphanumeric characters, _ or -”

None

CDPD-69279: Quick search does not return entities when using Japanese or Chinese characters to search properties

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Entities, such as `hdfs_path`, are not returned by the search when their indexable properties are searched using partial search terms made of Japanese or Chinese characters. Only exact matches return results when searching the indexed properties made of Japanese or Chinese characters.

None

CDPD-56085: [Impala Iceberg] LOAD DATA INPATH to Iceberg_table creates a temporary hive_table with name <iceberg_table_name>_tmp* and then marks it as DELETED in Atlas

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Running a query like `LOAD DATA INPATH` to `iceberg_table`, creates a temporary hive table with name `<iceberg_table_name>_tmp*` and then marks it as `DELETED` in Atlas. So in Atlas, a deleted entity is created corresponding to the temporary table `<iceberg_table_name>_tmp*`.

Tag added to the File system (HDFS) entity will not be propagated to the Iceberg table, user has to manually add to the Iceberg table, since the tag propagation is broken due to the deleted table in the flow.

CDPD-55301: The ddlQueries and ALTERTABLE_* lineage are missing for Spark tables created using spark3-shell

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

The `ddlQueries` and `outputFromProcesses` (lineage) is missing for the alter queries.

Apache JIRA: [ATLAS-4747](#)

CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeupException: null" is seen

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

CDPD-76035: Resource lookup for Atlas service is failing

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Once the Atlas configuration snippet `atlas.authentication.method.file` is enabled and a classification is created, these do not synchronize correctly to the Type Category resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the Type Name.

CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When creating tables with multiple levels of depth, a tag applied to the first table will be propagated along the lineage. After exporting the last table, it will not have the propagated tag.

Apache JIRA: [ATLAS-4889](#)

CDPD-77128: Log out message for Atlas needs to be more informative on Knox enabled cluster

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Logging out of Atlas does not manage the external authentication. Atlas does not communicate a log-out event with the external authentication management, Apache Knox. When you log out of Atlas, you can still open the instance of Atlas from the same web browser without re-authentication.

To prevent additional access to Atlas, close all browser windows and exit the browser.

CDPD-83398: REMOVE_LEGACY_REF_ATTRIBUTES Patch Skipped During Startup, Causing Errors While Processing DML Events

7.3.2 and its SPs and CHF, 7.3.1 and its SPs and CHF

The REMOVE_LEGACY_REF_ATTRIBUTES patch is skipped during Atlas patch processing, which leads to an error when Data Manipulation Language (DML) events (e.g., "insert data") are processed, events that are typically ignored. This is observed in environments where DML event filtering is not in place or the patch is not applied correctly. As a result, Atlas attempts to process these events and encounters errors related to missing or outdated attribute handling.

CDPD-88230: Bulk import operations currently do not construct lineage as expected

7.3.2 and its SPs and CHF, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its SPs and CHF

Bulk import operations do not create lineage consistently, even when the entities are imported successfully.

CDPD-88302: Bulk import operations currently do not create audits for the imported entities

7.3.2 and its SPs and CHF, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its SPs and CHF

Bulk import operations do not create audits for the imported entities.

CDPD-7982: HBase bridge stops at HBase table with deleted column family

7.3.2 and its SPs and CHF, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its SPs and CHF

Bridge importing metadata from HBase fails when it encounters an HBase table for which a column family was previously dropped. The error indicates:

```
Metadata service API
  org.apache.atlas.AtlasClientV2$API_V2@58112bc4 failed
  with status 404 (Not Found)
  Response Body
  ({"errorCode":"ATLAS-404-00-007","errorMessage":"Invalid instance
  creation/updation parameters passed :
  hbase_column_family.table: mandatory attribute v
  alue missing in type
  hbase_column_family"})
```

None

Apache JIRA: [ATLAS-3551](#)

CDPD-6675: Irregular qualifiedName format for Azure storage

7.3.2 and its SPs and CHF, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its SPs and CHF

The qualifiedName for hdfs_path entities created from Azure blob locations (ABFS) doesn't have the clusterName appended to it as do hdfs_path entities in other location types.

None

CDPD-922: IsUnique relationship attribute not honored

7.3.2 and its SPs and CHF, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its SPs and CHF

The Atlas model includes the ability to ensure that an attribute can be set to a specific value in only one relationship entity across the cluster metadata. For example, if you wanted to add metadata tags to relationships that you wanted to make sure were unique in the system, you could design the relationship attribute with the property "IsUnique" equal true. However, in this release, the IsUnique attribute is not enforced.

None

CDPD-1664: Guest users are redirected incorrectly

7.3.2 and its SPs and CHF, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its SPs and CHF

Authenticated users logging in to Atlas are redirected to the Cloudera Knox-based login page. However, if a guest user (without Atlas privileges) attempts to log in to Atlas, the user is redirected instead to the Atlas login page.

To avoid this problem, open the Atlas Dashboard in a private or incognito browser window.

CDPD-1823: Queries with ? wildcard return unexpected results

7.3.2 and its SPs and CHF, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its SPs and CHF

DSL queries in Advanced Search return incorrect results when the query text includes a question mark (?) wildcard character. This problem occurs in environments where trusted proxy for Knox is enabled, which is always the case for Cloudera.

None

CDPD-1884: Free text search in Atlas is case sensitive

7.3.2 and its SPs and CHF, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its SPs and CHF

The free text search bar in the top of the screen allows you to search across entity types and through all text attributes for all entities. The search shows the top 5 results that match the search terms at any place in the text (*term* logic). It also shows suggestions that match the search terms that begin with the term (term* logic). However, in this release, the search results are case-sensitive.

If you don't see the results you expect, repeat the search changing the case of the search terms.

CDPD-1892: Ranking of top results in free-text search not intuitive

7.3.2 and its SPs and CHF, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its SPs and CHF

The Free-text search feature ranks results based on which attributes match the search criteria. The attribute ranking is evolving and therefore the choice of top results may not be intuitive in this release.

If you don't find what you need in the top 5 results, use the full results or refine the search.

CDPD-3125: Logging out of Atlas does not manage the external authentication

7.3.2 and its SPs and CHF, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its SPs and CHF

At this time, Atlas does not communicate a log-out event with the external authentication management, Apache Knox. When you log out of Atlas, you can still open the instance of Atlas from the same web browser without re-authentication.

To prevent access to Atlas after logging out, close all browser windows and exit the browser.

CDPD-3160: Hive lineage missing for INSERT OVERWRITE queries

7.3.2 and its SPs and CHF, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its SPs and CHF

Lineage is not generated for Hive INSERT OVERWRITE queries on partitioned tables. Lineage is generated as expected for CTAS queries from partitioned tables.

None

CDPD-3208: Table alias values are not found in search

7.3.2 and its SPs and CHF, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its SPs and CHF

When table names are changed, Atlas keeps the old name of the table in a list of aliases. These values are not included in the search index in this release, so after a table name is changed, searching on the old table name will not return the entity for the table.

None

CDPD-4545: Searches for Qualified Names with "@" doesn't fetch the correct results

7.3.2 and its SPs and CHF, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its SPs and CHF

When searching Atlas qualifiedName values that include an "@" character (@), Atlas does not return the expected results or generate appropriate search suggestions.

Consider leaving out the portion of the search string that includes the @ sign, using the wildcard character * instead.

CDPD-4762: Spark metadata order may affect lineage

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Atlas may record unexpected lineage relationships when metadata collection from the Spark Atlas Connector occurs out of sequence from metadata collection from HMS. For example, if an ALTER TABLE operation in Spark changing a table name and is reported to Atlas before HMS has processed the change, Atlas may not show the correct lineage relationships to the altered table.

None

CDPD-10574: Suggestion order doesn't match search weights

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

At this time, the order of search suggestions does not honor the search weight for attributes.

None

Apache JIRA: [ATLAS-3081](#)

CDPD-10576: Deleted Business Metadata attributes appear in Search Suggestions

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Atlas search suggestions continue to show Business Metadata attributes even if the attributes have been deleted.

None

Apache JIRA: [ATLAS-3739](#)

CDPD-11692: Navigator table creation time not converted to Atlas

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

In converting content from Navigator to Atlas, the create time for Hive tables is not moved to Atlas.

None

Apache JIRA: [ATLAS-3739](#)

CDPD-11940: Database audit record misses table delete

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When a hive_table entity is created, the Atlas audit list for the parent database includes an update audit. However, at this time, the database does not show an audit when the table is deleted.

None

CDPD-11941: Table creation events missed when multiple tables are created in the same Hive command

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When multiple Hive tables are created in the same database in a single command, the Atlas audit log for the database may not capture all the table creation events. When there is a delay between creation commands, audits are created as expected.

None

CDPD-55122: Any user with ssh access can view the downloaded results

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs

When a user triggers a download from the Atlas user interface, the downloaded data can be viewed by any other user who has SSH access to the host. This occurs because the downloaded files are not sufficiently restricted by file permissions on the local filesystem, allowing unintended access to potentially sensitive data.

None

CDPD-67112: Import transforms do not work as expected when replacing a string which already has ":"

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs

The character ":" is not supported in path replacements. The import succeeds but location remains unchanged. The character " " must be avoided.

None

Apache JIRA: [ATLAS-4834](#)

CDPD-56590: Create table "like" from Iceberg table creates a hive_table instead of iceberg_table

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

By default, for tables created using the "like" command, lineage is not generated in Atlas. The destination like table should be of the same type as source table. Instead an iceberg_table for source and hive_table for destination are getting created.

CDPD-68191: Suggestions do not return the correct results when searching multiple Chinese characters

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Free text search does not return results when searching Chinese phrases made of multiple characters. Partial searches return the correct results.

None.

CDPD-77434: Search suggestions may fail to return results when using complete qualified names in RAZ-enabled environments

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When performing searches in the Atlas UI or through the REST API using a complete qualified name as the prefix string (for example, searching for an Iceberg column using its full qualified name), the suggestions API sometimes fails to return the expected entity in the suggestions list, even though the entity exists and would be returned in the actual search results. This issue occurs consistently in RAZ-enabled AWS clusters but has not been observed in IDB (internal data bus) or Cloudera on prem clusters. The behavior is intermittent: the REST API `/api/atlas/v2/search/suggestions?prefixString=<complete_qualifiedName>` may return an empty suggestions array (`{"suggestions": []}`), while the same entity appears correctly when searched through the user interface. This inconsistency affects user experience and makes entity discovery less reliable in RAZ-enabled environments.

None

CDPD-81499: Basic search performance degradation when filtering on boolean attributes

7.3.2 and its SPs and CHFs, 7.3.100 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When performing basic searches with filters on boolean attributes (for example, `isDisposed = true`), search operations may experience significant performance degradation, particularly with large datasets. This occurs because boolean datatype attributes are not indexed in Solr and require in-memory filtering in Atlas. When the search API attempts to return the requested number of results (for example, limit of 25), it repeatedly queries Solr for active entities and then filters them in Atlas based on boolean attribute values. This causes Atlas to traverse through potentially large numbers of entities until it finds enough matches, resulting in increased search times as data volume and query offset increase. Atlas logs may display warnings such as: `not using index-search for attribute EntityType.booleanAttribute; might cause poor performance.`

None

CDPD-69317: Export/Import API: When tag attributes types get modified between the exports, causes issues in import

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When using the Export/Import API, if a tag's attribute type is modified (for example, from string to double) between incremental exports, the import operation on the target cluster fails due to the difference in attribute types.

None

CDPD-65906: Export: When relationship is updated to block tag , the imported lineage doesn't have the updated relationship

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When using the Export/Import API to migrate data between clusters, if a relationship between entities is modified to block tag propagation (for example, setting propagation from table to process as 'None') and a tag is subsequently added, the tag correctly does not propagate on the source cluster. However, after exporting to the target cluster, the relationship on the target cluster does not reflect the blocked propagation, and the tag incorrectly propagates to the related entities.

None

Apache JIRA: [ATLAS-4833](#)

Known Issues in Apache Avro

Learn about the known issues in Avro, the impact or changes to the functionality, and the workaround.

Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known Issues in Cloud Connectors

Learn about the known issues in Cloud Connectors, the impact or changes to the functionality, and the workaround.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

There are no known issues carried forward that affect Cloud Connectors in Cloudera Runtime 7.3.2.

Known Issues in Cruise Control

Known issues and technical limitations for Cruise Control are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-44676: Rebalancing with Cruise Control does not work if the metric reporter fails to report the CPU usage metric

7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

If the CPU usage metric is not reported, the numValidWindows in Cruise Control will be 0 and proposal generation as well as partition rebalancing will not work. If this issue is present, the following message will be included in the Kafka logs:

```
WARN com.linkedin.kafka.cruisecontrol.metricsreporter.CruiseControlMetricsReporter:
    [CruiseControlMetricsReporterRunner]: Failed reporting CPU
    util.
```

```
java.io.IOException: Java Virtual Machine recent CPU usage is not
    available.
```

This issue is only known to affect Kafka broker hosts that have the following specifications:

- CPU: Intel(R) Xeon(R) CPU E5-2699 v4 @ 2.20GHz
- OS: Linux 4.18.5-1.el7.elrepo.x86_64 #1 SMP Fri Aug 24 11:35:05 EDT 2018 x86_64
- Java version: 8-18

Move the broker to a different machine where the CPU is different. This can be done by moving the host to a different cluster. For more information, see [Moving a Host Between Clusters](#)



Note: Cluster nodes affected by this issue are not displayed as unhealthy.

Known Issues in Apache HBase

Learn about the known issues in HBase, the impact or changes to the functionality, and the workaround.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

There are no known issues carried forward that affect HBase in Cloudera Runtime 7.3.2.

Known Issues in HDFS

Learn about the known issues in HDFS, the impact or changes to the functionality, and the workaround.

Known Issues Identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known Issues Identified Before Cloudera Runtime 7.3.2

There are no known issues carried forward that affect HDFS in Cloudera Runtime 7.3.2.

Known Issues in Apache Hive

Known issues and technical limitations for Hive are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2

CDPD-73781: High operational costs from automatic partition discovery on cloud storage

7.3.1 and its SPs and CHF, 7.2.18 and its SPs

In cloud environments, the Partition Management Task (PMT) could significantly increase operational costs and performance overhead for large tables with many partitions.

To address this issue, you can perform the following:

- Increase the value of the `metastore.partition.management.task.frequency` property to 86400 seconds.
- Disable partition discovery manually for specific tables incurring high costs by using the `discover.partitions` table property.

CDPD-93090: Failure when inserting data into partitions located on a different filesystem than the table

7.3.2

Inserting data into a table partition results in a `MoveTask` execution error if the partition is stored in a different filesystem than the parent table. For example, if you create an external table in S3 but add a partition with an HDFS location, the `INSERT` operation fails with a `Wrong FS` error. Although the data is physically inserted into the correct location, the query returns a non-zero exit code and displays an error message.

None

CDPD-92560: Partition insertion errors across multiple buckets in RAZ enabled clusters

7.3.2

Inserting data into a table partition results in a `MoveTask` execution error when the partition is located in a different S3 bucket than the base table. In Ranger Access Control (RAZ) enabled Private Cloud clusters, the `INSERT` operation fails with a `Wrong FS` error, even though the data is physically written to the specified partition location.

None

CDPD-74192: Hive SSL certificate error during large record copies

7.3.1 and its SPs and CHF, 7.2.18.300, 7.2.18.400

When you copy a large number of records (approximately 1 million or more) using Hive, the operation might fail with a

```
javax.net.ssl.SSLException: org.bouncycastle.tls.TlsFatalAlert:
certificate_unknown(46)
```

error. This issue occurs during the job commit phase when Hive attempts to execute HTTP requests to S3 for file operations such as `copyFile`.

None

CDPD-68096: Unable to set S3 credentials at the session level in Hive

7.3.2

Accessing S3 buckets from Hive on on premises results in a `MetaException` or `AccessDeniedException` when you attempt to create external tables by using session-level credentials. While Spark and other Hadoop jobs allow you to define `fs.s3a.access.key` and `fs.s3a.secret.key` at the session level, Hive Metastore (HMS) does not honor these session-specific parameters during path validation.

To access S3 buckets, you must configure the S3 credentials at the global level in the Hive or Hadoop cluster configuration (such as `core-site.xml`) or use a supported credential provider like the Hadoop Credential Shell to store keys securely. Setting these parameters at the session level is currently not supported for Hive table creation.

[HIVE-16913](#)

Known issues identified before Cloudera Runtime 7.3.2**DWX-22436: DL upgrade recovery fails due to Metastore schema incompatibility**

7.3.1.600

When attempting a Data Lake (DL) upgrade recovery from version 7.2.18.1100 to Cloudera Runtime 7.3.1.500, the process fails because the Hive Metastore schema versions are incompatible. The error indicates a mismatch between the Hive version (3.1.3000.7.3.1.500-182) and the database schema version (3.1.3000.7.2.18.0-Update2). This blocks Data Lake recovery if an upgrade fails, impacting customers.

Before you initiate the recovery process, manually update the Hive Metastore schema to match the target version by using the `schematool` utility.

1. Obtain the Hive database password: Run the following command to retrieve the password from the pillar configuration:

```
cat /srv/pillar/postgresql/postgre.sls
```

2. Back up the existing configuration: Move the current configuration directory to a backup location:

```
mv /etc/hive/conf /etc/hive/conf_backup
mkdir /etc/hive/conf
```

3. Prepare the temporary configuration: Copy the process files to the new configuration directory:

```
scp /var/run/cloudera-scm-agent/process/<process-id>-hive-metastore-create-tables/* /etc/hive/conf/
```

4. Update the connection password: Open the `/etc/hive/conf/hive-site.xml` file and perform the following modifications:

- Set the `javax.jdo.option.ConnectionPassword` property to your Hive database password.
- Comment out the `hadoop.security.credential.provider.path` property.

5. Run the schema upgrade tool: Execute the `schematool` to synchronize the version:

```
/opt/cloudera/parcels/CDH/lib/hive/bin/schematool -dbType postgres -initOrUpgradeSchema --verbose
```

6. Restore the original configuration: Remove the temporary directory and restore your backup:

```
rm -rf /etc/hive/conf
mv /etc/hive/conf_backup /etc/hive/conf
```

7. Restart the cluster: Restart the services to initialize the Hive Metastore with the updated schema.

CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

7.1.9 SP1 CHF4, 7.1.7 SP3 CHF7, 7.3.1.100, and its higher versions

Atlas hook authorization error causes `HiveCreateSysDb` command to time out due to repeated retries.

None

CDPD-74680: DAG not retried after failure

7.3.1 and its higher versions

When executing a Hive query, if the `ApplicationMaster` container fails, Hive does not retry the DAG if the failure message contains some diagnostic information including a line break, leading to query failure (instead of retry).

None

Known Issues in Cloudera Data Explorer (Hue)

Known issues and technical limitations for Cloudera Data Explorer (Hue) are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2

CDPD-98754: Data Explorer cannot access HDFS in Data Discovery Template

7.3.2

In the Data Discovery Template, Data Explorer cannot access HDFS. This issue occurs because Cloudera Manager assigns a standby HDFS NameNode to the `webhdfs_url` property instead of the active NameNode, which prevents the file browser from displaying HDFS data.

1. Log in to Cloudera Manager as an Administrator.
2. Go to `Clusters HDFS` to identify which NameNode is currently active.
3. Go to `Clusters Data Explorer (Hue) Configuration`.
4. Change the value of the `webhdfs_url` property from the standby NameNode to the active NameNode.
5. Click `Save Changes`.
6. Restart the Data Explorer service.

CDPD-95911: File operations fail for files with special characters in GS RAZ environments

7.3.2

In GS RAZ environments, you cannot rename, move, or copy files that contain special characters such as the percent sign (%). When you attempt these operations, the `Failed to receive a valid response from RAZ` error message is displayed. Additionally, operations on files containing the hash character (#) do not result in an error, but the action is not performed.

None.

CDPD-94944: File and directory names contain special characters after a rename operation

7.3.2

In Data Explorer, file and directory names that contain special characters, such as the ampersand (&) and the apostrophe ('), are incorrectly transformed into HTML entities after a rename operation. For example, the ampersand is replaced by `&` and the apostrophe is replaced by `'`. This issue occurs during rename operations for files with names containing characters such as `~@$&()*!+'=;.txt`.

None.

CDPD-95666: Endless progress bar for S3 folders without read permissions on Azure

7.3.2

In Azure environments, a persistent loading bar is displayed when you select a folder or file for which you do not have read permissions. This issue occurs across the file system, including S3 and HDFS. Currently, the user interface does not display an alert message to indicate that permissions are missing. Instead, the UI remains in a loading state.

None.

CDPD-93181: Incorrect Data Explorer server status reporting

Data Explorer server status is incorrectly reported as `Green (Healthy)` in Cloudera Manager even if the process fails to start. This occurs when the `mysqlclient` dependency is missing, causing database migrations to fail. The Data Explorer server appears healthy because the `loglistener.py` sidecar process starts first and maintains a Process ID (PID). Cloudera Manager monitors this PID and reports a successful start, even though the main Data Explorer process has terminated due to the missing dependency.

1. Stop the Hue service in Cloudera Manager.
2. Log in to the affected Hue host and manually terminate the loglistener.py sidecar process.
3. Remove any stale PID files located in the /tmp/ directory, for example, /tmp/hue_*.pid.
4. Ensure that the mysqlclient dependency is installed on the host before attempting to restart the service.

Known issues identified Before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-95086: Data Explorer UI fails to load due to a connection timeout error

7.3.2, 7.3.1.706

In a Cloudera Data Hub high availability (HA) environment, the Data Explorer UI can fail to load after the master0 HDFS node is deleted. Data Explorer continues to attempt connections to the deleted node because Data Explorer does not support HDFS HA. As a result, the Data Explorer file browser points to an invalid HDFS endpoint, which causes connection timeouts and renders the UI inaccessible.

You can restore access to the Data Explorer UI by updating the configured HttpFS node to an active HDFS node in Cloudera Manager and then restarting the Data Explorer service.

CDPD-92946: Data Explorer debug logging cannot be disabled

7.3.2, 7.3.1.600, 7.3.1.706

Data Explorer debug logging is always enabled and cannot be turned off.

None.

CDPD-88964: Data Explorer logs missing in Data Explorer UI

7.3.2, 7.3.1.500 through 7.3.1.706

The Data Explorer UI might display only a few lines of logs instead of the complete Data Explorer logs. This can occur due to leftover Gunicorn processes that interfere with the proper logging and display of logs within the Data Explorer interface.

1. Stop the Data Explorer service.
2. Terminate any remaining Gunicorn processes to clear hung or orphan processes that might be causing the issue. Run the following command (use sudo if not running as root):

```
# pids=$(ps -efwww | grep rungunicornserver | grep -v "grep" |
grep "rungunicornserver"| awk '{ print $2 }') && for i in $pi
ds; do kill -9 $i ; done
```

3. Restart the Data Explorer service.

CDPD-90510: Defunct Data Explorer Gunicorn worker processes accumulate

7.3.2, 7.3.1.500 through 7.3.1.706

Data Explorer on Ubuntu 22 using Oracle Database can accumulate defunct rungunicornserver worker processes due to incomplete process termination and stale database connections. This can lead to a cluttered process table, which does not critically impact service functionality.

Periodically clean defunct worker processes by running the following command:

```
pgrep -f 'hue rungunicornserver' | xargs -r kill -9
```

OPSAPS-75134: LDAP and Kerberos dual authentication fails with HiveOnTez in HTTP Transport Mode

7.3.2, 7.3.1 and its SPs and CHF, Cloudera Manager 7.13.1 and its SPs and CHF

Enabling LDAP for the HiveOnTez service in a Kerberos environment with transport mode set to HTTP, Data Explorer fails to load database information . This failure occurs because Data Explorer does not support the combined KERBEROS and LDAP authentication value required by HiveOnTez. The resulting conflict prevents a successful connection to HiveServer2

- Log in to Cloudera Manager Hive-On-Tez Configuration Hive Service Advanced Configuration Snippet (Safety Valve) for hive-site.xml and set the following value:

```
hive.server2.authentication=LDAP , KERBEROS
```

- Go to Hue Configuration Hue Server Advanced Configuration Snippet (Safety Valve) for hive-site.xml and set the following value:

```
hive.server2.authentication=KERBEROS
```

CDPD-58978: Batch query execution using Data Explorer fails with Kerberos error

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

When you run Impala queries in a batch mode, you encounter failures with a Kerberos error even if the keytab is configured correctly. This is because submitting Impala, Sqoop, Pig, or pyspark queries in a batch mode launches a shell script Oozie job from Data Explorer and this is not supported on a secure cluster.

None. Submit the queries individually.

CDPD-59677: Unable to view Phoenix tables on the left assist in Data Explorer

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

On clusters secured with Knox, you might not be able to see Phoenix tables on the left assist that are present under the default database (that is, an empty("") database).

None.

Known Issues in Apache Iceberg

Learn about the known issues in Iceberg, the impact or changes to the functionality, and the workaround.

Known Issues identified in Cloudera Runtime 7.3.2

DWX-18843: Unable to read Iceberg table from Hive Virtual Warehouse

7.3.2

If you have used Apache Flink to insert data into an Iceberg table that is created from Hive, you cannot read the Iceberg table from the Hive Virtual Warehouse.

Add the engine.hive.enabled table property through the Hive beeline and set the value to "true". You can add this table property either while creating the Iceberg table or use the ALTER TABLE statement to add the table property.

DWX-18489: Hive compaction of Iceberg tables results in a failure

7.3.2

When Cloudera Data Warehouse and Cloudera Data Hub are deployed in the same environment and use the same Hive Metastore (HMS) instance, the Cloudera Data Hub compaction workers can inadvertently pick up Iceberg compaction tasks. Since Iceberg compaction is not yet supported in the latest Cloudera Data Hub version, the compaction tasks will fail when they are processed by the Cloudera Data Hub compaction workers.

In such a scenario where both Cloudera Data Warehouse and Cloudera Data Hub share the same HMS instance and there is a requirement to run both Hive ACID and Iceberg compaction jobs, it is recommended that you use the Cloudera Data Warehouse environment for these jobs. If you want to

run only Hive ACID compaction tasks, you can choose to use either the Cloudera Data Warehouse or Cloudera Data Hub environments.

If you want to run the compaction jobs without changing the environment, it is recommended that you use Cloudera Data Warehouse. To avoid interference from Cloudera Data Hub, change the value of the `hive.compactor.worker.threads` Hive Server (HS2) property to '0'. This ensures that the compaction jobs are not processed by Cloudera Data Hub.

1. In Cloudera Manager, click `Clusters Hive Configuration` to navigate to the configuration page for HMS.
2. Search for `hive.compactor.worker.threads` and modify the value to '0'.
3. Save the changes and restart the Hive service.

DWX-17254: Merging Iceberg branches requires a target table alias

7.3.2

Hive supports only one level of qualifier when referencing columns. In other words only one dot is accepted. For example, `select table.col from ...;` is allowed. `select db.table.col` is not allowed. Using the merge statement to merge Iceberg branches without a target or source table alias causes an exception:

```
org.apache.hadoop.hive.ql.parse.SemanticException: ... Invalid table alias or column reference ...
```

Use an alias, for example `t`, for the target table.

```
merge into mydb.target.branch_branch1 t using mydb.source.branch_branch1 s on t.id = s.id when matched then update set value = 'matched';
```

Apache Jira: [HIVE-28055](#)

DWX-17210, DWX-13733: Timeout issue querying Iceberg tables from Hive

7.3.2

When querying Iceberg tables from Hive, the queries can fail due to a timeout issue.

1. Add the following configurations to `hadoop-core-site` for the Database Catalog and the Virtual Warehouse.
 - `fs.s3.maxConnections=1000`
 - `fs.s3a.connection.maximum=1000`
2. Restart the Database Catalog and Virtual Warehouse.

DWX-14163: Limitations reading Iceberg tables in Avro file format from Impala

7.3.2

The Avro, Impala, and Iceberg specifications describe some limitations related to Avro, and those limitations exist in Cloudera. In addition to these, the DECIMAL type is not supported in this release.

None.

DEX-7946: Data loss during migration of a Hive table to Iceberg

7.3.2

In this release, by default the table property `'external.table.purge'` is set to true, which deletes the table data and metadata if you drop the table during migration from Hive to Iceberg.

Either one of the following workarounds prevents data loss during table migration:

- Set the table property `'external.table.purge'='FALSE'`.
- Do not drop a table during migration from Hive to Iceberg.

DWX-13062: Converting a Hive table having CHAR or VARCHAR columns to Iceberg causes an exception

7.3.2

CHAR and VARCHAR data can be shorter than the length specified by the data type. Remaining characters are padded with spaces. Data is converted to a string in Iceberg. This process can yield incorrect results when you query the converted Iceberg table.

Change columns from CHAR or VARCHAR to string types before converting the Hive table to Iceberg.

Apache Jira: [HIVE-26507](#)

Known issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-92182: Inserting into Hive Iceberg tables on S3 fails with RazS3ClientCredentialsException

7.3.2, 7.3.1.706, 7.3.1.600, 7.3.1.500

In RAZ-enabled clusters where HDFS is the default file system, attempts to insert data into Hive Iceberg tables pointing explicitly to S3 locations fail with a RazS3ClientCredentialsException

None.

CDPD-89390/CDPD-83022: Incorrect row count displayed in table metadata after compaction

7.3.2, 7.3.1.600, 7.3.1.500, 7.3.1.400

After running data compaction operations on large tables, the row count displayed by the DESC RIBE FORMATTED command may be inaccurate. Initially, the count may appear higher than the actual number of rows. Subsequently, after running the ANALYZE command to update table statistics, the count might then appear lower than the actual number of rows.

This issue has been observed in large tables containing a significant number of historical snapshots (exceeding 10000). All these snapshots are primarily generated through UPDATE operations.

It is important to note that this is just a metadata display issue, and there is no loss of data. The underlying table data remains complete and correct.

To obtain an accurate row count, use the SELECT COUNT(*) query.

Known Issues in Apache Impala

Known issues and technical limitations for Impala are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2**CDPD-90807: Thrift protocol limitation during Impala zero downtime upgrade (ZDU)**

7.3.1.500 through 7.3.1.706, 7.3.2

Zero Downtime Upgrades (ZDU) for Impala are not supported when upgrading to version 7.3.2.

This is due to Thrift protocol incompatibilities that can cause queries to fail during upgrade.

None

CDPD-90250: Incorrect file storage location for Impala tables when using S3 as default filesystem

7.3.2

When you created an external table in Impala by using the LOCATION parameter, the files were stored in S3 instead of HDFS. This occurred even if the provided path was intended for HDFS, provided that S3 was configured as the default filesystem (fs.defaultFS) in the cluster.

This inconsistency leads to `AccessDeniedException` errors if the Impala service does not have the necessary Ranger permissions to write to the S3 bucket.

None

IMPALA-14472: Writing arrays to Kudu tables

7.3.2

Impala does not currently support writing array data into Kudu tables.

None

Apache Jira: [IMPALA-14472](#)

Known issues identified before Cloudera Runtime 7.3.2

DWX-20490: Impala queries fail with "Caught exception The read operation timed out, type=<class 'socket.timeout'> in ExecuteStatement"

7.3.1.500

Queries in `impala-shell` fail with a socket timeout error in `execute statement` which submits the query to the coordinator. The error occurs when query execution takes longer to start mainly when query planning is slow due to frequent metadata changes.

Increase the socket timeout on the client side. Set `--client_connect_timeout_ms` to a higher value, e.g. add `--client_connect_timeout_ms=600000` to the `impala-shell` command line.

DWX-20491: Impala queries fail with EOFException: End of file reached before reading fully

7.3.1.500

Impala queries fail with an `EOFException` when reading from an HDFS file stored in an S3A location. The error occurs when the file is removed. If the file is removed using SQL commands like `DROP PARTITION`, there may be a significant lag in Hive Metastore event processing. If removed by non-SQL operations, run `REFRESH` or `INVALIDATE METADATA` on the table to resolve the issue.

Run `REFRESH/INVALIDATE METADATA <table>;`

CDPD-94720: Impala startup failure due to invalid TLS v1.3 ciphers

7.3.1 and its higher versions

When running Impala on a machine with OpenSSL 1.1.1, providing an invalid or TLS v1.2 ciphersuite in the `--tls_ciphersuites` startup flag causes the process to fail during startup. While OpenSSL 3.x ignores invalid ciphers, OpenSSL 1.1.1 returns an error if any ciphersuite in the list is invalid, even if other valid TLS v1.3 ciphers are present.

Ensure that the list in the `--tls_ciphersuites` startup flag contains only valid TLS v1.3 ciphersuites and does not contain any TLS v1.2 ciphersuites.

Apache Jira: [IMPALA-14625](#)

IMPALA-532: Impala should tolerate bad locale settings

7.3.1 and its higher versions

If the `LC_*` environment variables specify an unsupported locale, Impala does not start.

Add `LC_ALL="C"` to the environment settings for both the Impala daemon and the Statestore daemon.

IMPALA-691: Process mem limit does not account for the JVM's memory usage

7.3.1 and its higher versions

Some memory allocated by the JVM used internally by Impala is not counted against the memory limit for the `impalad` daemon.

To monitor overall memory usage, use the `top` command, or add the memory figures in the Impala web UI `/memz` tab to JVM memory usage shown on the `/metrics` tab.

IMPALA-635: Avro Scanner fails to parse some schemas

7.3.1 and its higher versions

The default value in Avro schema must match type of first union type, e.g. if the default value is null, then the first type in the UNION must be "null".

Swap the order of the fields in the schema specification. For example, use ["null", "string"] instead of ["string", "null"]. Note that the files written with the problematic schema must be rewritten with the new schema because Avro files have embedded schemas.

IMPALA-1024: Impala BE cannot parse Avro schema that contains a trailing semi-colon

7.3.1 and its higher versions

If an Avro table has a schema definition with a trailing semicolon, Impala encounters an error when the table is queried.

Remove trailing semicolon from the Avro schema.

IMPALA-1652: Incorrect results with basic predicate on CHAR typed column

7.3.1 and its higher versions

When comparing a CHAR column value to a string literal, the literal value is not blank-padded and so the comparison might fail when it should match.

Use the RPAD() function to blank-pad literals compared with CHAR columns to the expected length.

IMPALA-1821: Casting scenarios with invalid/inconsistent results

7.3.1 and its higher versions

Using a CAST() function to convert large literal values to smaller types, or to convert special values such as NaN or Inf, produces values not consistent with other database systems. This could lead to unexpected results from queries.

None

IMPALA-2005: A failed CTAS does not drop the table if the insert fails

7.3.1 and its higher versions

If a CREATE TABLE AS SELECT operation successfully creates the target table but an error occurs while querying the source table or copying the data, the new table is left behind rather than being dropped.

Drop the new table manually after a failed CREATE TABLE AS SELECT

IMPALA-3509: Breakpad minidumps can be very large when the thread count is high

7.3.1 and its higher versions

The size of the breakpad minidump files grows linearly with the number of threads. By default, each thread adds 8 KB to the minidump size. Minidump files could consume significant disk space when the daemons have a high number of threads.

Add `-\-minidump_size_limit_hint_kb=size` to set a soft upper limit on the size of each minidump file. If the minidump file would exceed that limit, Impala reduces the amount of information for each thread from 8 KB to 2 KB. (Full thread information is captured for the first 20 threads, then 2 KB per thread after that.) The minidump file can still grow larger than the "hinted" size. For example, if you have 10,000 threads, the minidump file can be more than 20 MB.

IMPALA-4978: Impala requires FQDN from hostname command on Kerberized clusters

7.3.1 and its higher versions

The method Impala uses to retrieve the host name while constructing the Kerberos principal is the `gethostname()` system call. This function might not always return the fully qualified domain name, depending on the network configuration. If the daemons cannot determine the FQDN, Impala does not start on a Kerberized cluster.

Test if a host is affected by checking whether the output of the `hostname` command includes the FQDN. On hosts where `hostname` only returns the short name, pass the command-line flag `##hostname=FULLY_QUALIFIED_DOMAIN_NAME` in the startup options of all Impala-related daemons.

IMPALA-6671: Metadata operations block read-only operations on unrelated tables

7.3.1 and its higher versions

Metadata operations that change the state of a table, like `COMPUTE STATS` or `ALTER RECOVER PARTITIONS`, may delay metadata propagation of unrelated unloaded tables triggered by statements like `DESCRIBE` or `SELECT` queries.

None

IMPALA-7072: Impala does not support Heimdal Kerberos

None

CDPD-28139: Set spark.hadoop.hive.stats.autogather to false by default

As an Impala user, if you submit a query against a table containing data ingested using Spark and you are concerned about the quality of the query plan, you must run `COMPUTE STATS` against such a table in any case after an ETL operation because `numRows` created by Spark could be incorrect. Also, use other stats computed by `COMPUTE STATS`, e.g., Number of Distinct Values (NDV) and NULL count for good selectivity estimates.

For example, when a user ingests data from a file into a partition of an existing table using Spark, if `spark.hadoop.hive.stats.autogather` is not set to `false` explicitly, `numRows` associated with this partition would be 0 even though there is at least one row in the file. To avoid this, the workaround is to set `"spark.hadoop.hive.stats.autogather=false"` in the "Spark Client Advanced Configuration Snippet (Safety Valve) for `spark-conf/spark-defaults.conf`" in Spark's CM Configuration section.

IMPALA-2422: % escaping does not work correctly when occurs at the end in a LIKE clause

7.3.1 and its higher versions

If the final character in the RHS argument of a `LIKE` operator is an escaped `\%` character, it does not match a `%` final character of the LHS argument.

None

IMPALA-2603: Crash: impala::Coordinator::ValidateCollectionSlots

A query could encounter a serious error if includes multiple nested levels of `INNER JOIN` clauses involving subqueries.

None

CDPD-59625: Impala shell in RHEL 9 with Python 2 as default does not work

7.1.9, 7.3.1 and its higher version

If you try to run `impala-shell` on RHEL 9 by setting the default python executable available in `PATH` to Python 2, it will fail since RHEL 9 is compatible only with Python 3.

If you run into such issues, set this parameter pointing to Python 3, `IMPALA_PYTHON_EXECUTABLE=python3`.

Impala cannot update table if the 'external.table.purge' property is not set to true

Impala cannot update a table using DDL statements if the `'external.table.purge'` property is `FALSE`. `ALTER TABLE` statements return success with no changes to the table.

`ALTER TABLE` statements should be issued twice if `"external.table.purge"` was `FALSE` initially.



Note: For Iceberg tables, Impala users should always use `CREATE TABLE` since the Impala `CREATE TABLE` statement sets the flag to `TRUE`. However, Impala `CREATE EXTERNAL TABLE` sets the flag to `FALSE`.

Impala's known limitation when querying compacted tables

When the compaction process deletes the files for a table from the underlying HDFS location, the Impala service does not detect the changes as the compactions does not allocate new write ids. When the same table is queried from Impala it throws a 'File does not exist' exception that looks something like this:

```
Query Status: Disk I/O error on <node>:22000: Failed to open HDFS file hdfs://nameservice1/warehouse/tablespace/managed/hive/<database>/<table>/xxxxx
Error(2): No such file or directory Root cause: RemoteException: File does not exist: /warehouse/tablespace/managed/hive/<database>/<table>/xxxx
```

Use the [REFRESH/INVALIDATE](#) statements on the affected table to overcome the 'File does not exist' exception.

Impala api calls via Knox require configuration if the Knox customized Kerberos principal name is a default service user name

To access Impala API calls via Knox, if the Knox customized Kerberos principal name is a default service user name, then configure "authorized_proxy_user_config" by clicking Clusters->Impala->configuration. Include the Knox customized Kerberos principal name in the comma separated list of values <knox_custom_kerberos_principal_name>="*" where <knox_custom_kerberos_principal_name> is the value of the Kerberos Principal in the Knox service. Select Clusters>Knox>Configuration and search for Kerberos Principal to display this value.

CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes.

7.1.7

You must use a single Knox node to access Impala UI.

CDPD-21828: Multiple permission assignment through grant is not working

7.1.7

None

IMPALA-11871: INSERT statement does not respect Ranger policies for HDFS

7.3.1, 7.3.1.300 and its higher version

7.3.1.100, 7.3.1.200

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to cm_hdfs -> all-path for the user impala, inserting into a table whose HDFS POSIX permissions happen to exclude impala access will result in "AnalysisException: Unable to INSERT into target table (default.t1) because Impala does not have WRITE access to HDFS location: hdfs://XXXXXXXXXXXXX"

OPSAPS-46641: A single parameter exists in Cloudera Manager for specifying the Impala Daemon Load Balancer. Because BDR and Hue need to use different ports when connecting to the load balancer, it is not possible to configure the load balancer value so that BDR and Hue will work correctly in the same cluster.

The workaround is to use the load balancer configuration either without a port specification, or with the Beeswax port: this will configure BDR. To configure Hue use the "Hue Server Advanced Configuration Snippet (Safety Valve) for impalad_flags" to specify the the load balancer address with the HiveServer2 port.

Impala known limitation when querying compacted tables

7.3.1 and its higher versions

When the compaction process deletes the files for a table from the underlying HDFS location, the Impala service does not detect the changes as the compactions does not allocate new write ids.

When the same table is queried from Impala it throws a 'File does not exist' exception that looks something like this:

```
Query Status: Disk I/O error on <node>:22000: Failed to open HDFS file hdfs://nameservice1/warehouse/tablespace/managed/hive/<database>/<table>/xxxxx
Error(2): No such file or directory Root cause: RemoteException:
  File does not exist: /warehouse/tablespace/managed/hive/<database>/<table>/xxxx
```

Use the [REFRESH/INVALIDATE](#) statements on the affected table to overcome the 'File does not exist' exception.

Impala Virtual Warehouses might produce an error when querying transactional (ACID) tables

Problem: If you are querying transactional (ACID) tables with an Impala Virtual Warehouse and compaction is run on the compacting Hive Virtual Warehouse, the query might fail. The compacting process deletes files and the Impala Virtual Warehouse might not be aware of the deletion. Then when the Impala Virtual Warehouse attempts to read the deleted file, an error can occur. This situation occurs randomly.

Run the INVALIDATE METADATA statement on the transactional (ACID) table to refresh the metadata. This fixes the problem until the next compaction occurs.

IMPALA-5605: Configuration to prevent crashes caused by thread resource limits

Impala could encounter a serious error due to resource usage under very high concurrency. The error message is similar to:

```
F0629 08:20:02.956413 29088 llvm-codegen.cc:111] LLVM hit fatal
error: Unable to allocate section memory!
terminate called after throwing an instance of 'boost::exception_
detail::clone_impl<boost::exception_detail::error_info_injector<
boost::thread_resource_error> >'
```

To prevent such errors, configure each host running an `impalad` daemon with the following settings:

```
echo 2000000 > /proc/sys/kernel/threads-max
echo 2000000 > /proc/sys/kernel/pid_max
echo 8000000 > /proc/sys/vm/max_map_count
```

Add the following lines in `/etc/security/limits.conf`:

```
impala soft nproc 262144
impala hard nproc 262144
```

IMPALA-9350: Ranger audit logs for applying column masking policies missing

Impala is not producing these logs.

None

IMPALA-1792: ImpalaODBC: Can not get the value in the SQLGetData(m-x th column) after the SQLBindCol(m th column)

If the ODBC `SQLGetData` is called on a series of columns, the function calls must follow the same order as the columns. For example, if data is fetched from column 2 then column 1, the `SQLGetData` call for column 1 returns NULL.

Fetch columns in the same order they are defined in the table.

Known Issues in Apache Kafka

Known issues and technical limitations for Kafka are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

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

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

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

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

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

KAFKA-2561: Performance degradation when SSL Is enabled

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

In some configuration scenarios, significant performance degradation can occur when SSL is enabled. The impact varies depending on your CPU, JVM version, Kafka configuration, and message size. Consumers are typically more affected than producers.

Configure brokers and clients with `ssl.secure.random.implementation = SHA1PRNG`. It often reduces this degradation drastically, but its effect is CPU and JVM dependent.

RANGER-3809: Idempotent Kafka producer fails to initialize due to an authorization failure

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

Kafka producers that have idempotence enabled require the Idempotent Write permission to be set on the cluster resource in Ranger. If permission is not given, the client fails to initialize and an error similar to the following is thrown:

```
org.apache.kafka.common.KafkaException: Cannot execute transactional method because we are in an error state
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeFailWithError(TransactionManager.java:1125)
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeAddPartition(TransactionManager.java:442)
```

```

        at org.apache.kafka.clients.producer.KafkaProducer.doSend(KafkaProducer.java:1000)
        at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:914)
        at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:800)
        .
        .
        .
        Caused by: org.apache.kafka.common.errors.ClusterAuthorizationException: Cluster authorization failed.

```

Idempotence is enabled by default for clients in Kafka 3.0.1, 3.1.1, and any version after 3.1.1. This means that any client updated to 3.0.1, 3.1.1, or any version after 3.1.1 is affected by this issue.

This issue has two workarounds, do either of the following:

- Explicitly disable idempotence for the producers. This can be done by setting `enable.idempotence` to `false`.
- Update your policies in Ranger and ensure that producers have Idempotent Write permission on the cluster resource.

CDPD-49304: AvroConverter does not support composite default values

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2.0

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

CFM-3532: The Stateless NiFi Source, Stateless NiFi Sink, and HDFS Stateless Sink connectors cannot use Snappy compression

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

This issue only affects Stateless NiFi Source and Sink connectors if the connector is running a dataflow that uses a processor that uses Hadoop libraries and is configured to use Snappy compression. The HDFS Stateless Sink connector is only affected if the `Compression Codec` or `Compression Codec` for Parquet properties are set to `SNAPPY`.

If you are affected by this issue, errors similar to the following will be present in the logs.

```
Failed to write to HDFS due to java.lang.UnsatisfiedLinkError: org.apache.hadoop.util.NativeCodeLoader.buildSupportsSnappy()
```

```
Failed to write to HDFS due to java.lang.RuntimeException: native snappy library not available: this version of libhadoop was built without snappy support.
```

Download and deploy missing libraries.



Important: Ensure that you complete steps 1-11 on all Kafka Connect hosts. Additionally, ensure that the advanced configuration snippet in step 12 is configured for all Kafka Connect role instances.

1. Create the `/opt/nativelibs` directory.

```
mkdir /opt/nativelibs
```

2. Change the owner to kafka.

```
chown kafka:kafka /opt/nativelibs
```

3. Locate the directory containing the Hadoop native libraries and copy its contents to the directory you created.

```
cp /opt/cloudera/parcels/CDH/lib/hadoop/lib/native/* /opt/nativelibs
```

4. Verify that libsnappy.so was copied to the directory you created.
5. Remove the following from /opt/nativelibs.

```
libhadoop.a
libhadoop.so
libhadoop.so.1.0.0
```

6. Run the following command.

```
hadoop version
```

The command returns the Hadoop version running in the cluster. Note down the first three digits in the version.

7. Go to <https://archive.apache.org/dist/hadoop/common/> and download the Hadoop version that matches the first three digits of the version running in the cluster.

For example, if your Hadoop version is 3.1.1.7.1.9.0-296, then you need to download Hadoop 3.1.1.

8. Extract the downloaded archive.
9. Copy the following libraries from the downloaded archive to /opt/nativelibs on the cluster host.

```
libhadoop.a
libhadoop.so.1.0.0
```

The libraries are located in `hadoop-***VERSION***/lib/native`.

10. Create a symlink named libhadoop.so and point it to /opt/nativelibs/libhadoop.so.1.0.0.

```
ln -s /opt/nativelibs/libhadoop.so.1.0.0 /opt/nativelibs/libhadoop.so
```

11. Change the owner of every entry within /opt/nativelibs to kafka.

```
chown -h kafka:kafka /opt/nativelibs/*
```

12. In Cloudera Manager, go to **Kafka service Configuration**.
13. Add the following key-value pair to **Kafka Connect Environment Advanced Configuration Snippet (Safety Valve)**.
 - Key: LD_LIBRARY_PATH
 - Value: /opt/nativelibs
14. Click **Save Changes**.
15. Restart the Kafka service.

Unsupported Features

The following Kafka features are not supported in Cloudera:

- Only Java and .Net based clients are supported. Clients developed with C, C++, Python, and other languages are currently not supported.

- The Kafka default authorizer is not supported. This includes setting ACLs and all related APIs, broker functionality, and command-line tools.
- SASL/SCRAM is only supported for delegation token based authentication. It is not supported as a standalone authentication mechanism.

Limitations

Collection of partition level metrics may cause Cloudera Manager performance to degrade

If the Kafka service operates with a large number of partitions, collection of partition level metrics may cause Cloudera Manager performance to degrade.

If you are observing performance degradation and your cluster is operating with a high number of partitions, you can choose to disable the collection of partition level metrics.



Important: If you are using Streams Messaging Manager to monitor Kafka or Cruise Control for rebalancing Kafka partitions, be aware that both Streams Messaging Manager and Cruise Control rely on partition level metrics. If partition level metric collection is disabled, Streams Messaging Manager will not be able to display information about partitions. In addition, Cruise Control will not operate properly.

Complete the following steps to turn off the collection of partition level metrics:

1. Obtain the Kafka service name.
 - a. In Cloudera Manager, Select the Kafka service.
 - b. Select any available chart, and select Open in Chart Builder from the configuration icon drop-down.
 - c. Find \$SERVICENAME= near the top of the display.
The Kafka service name is the value of \$SERVICENAME.
2. Turn off the collection of partition level metrics.
 - a. Go to Hosts Hosts Configuration .
 - b. Find and configure the Cloudera Manager Agent Monitoring Advanced Configuration Snippet (Safety Valve) configuration property.

Enter the following to turn off the collection of partition level metrics:

```
[KAFKA_SERVICE_NAME]_feature_send_broker_topic_partition_entity_update_enabled=false
```

Replace [KAFKA_SERVICE_NAME] with the service name of Kafka obtained in step 1. The service name should always be in lower case.

- c. Click Save Changes.

Known Issues in Apache Knox

Learn about the known issues in Apache Knox, the impact or changes to the functionality, and the workaround.

Known Issues identified in Cloudera Runtime 7.3.2

OPSAPS-76713: Knox crashes under high concurrent load when using PAM authentication on RHEL 9 or Rocky Linux 9

7.3.2

On RHEL 9 or Rocky Linux 9 versions 9.1 through 9.7, a concurrency bug in libeconf 0.4.1 causes libpam to crash under high concurrent load. This affects Knox when the cdp-proxy-api topology uses PAM authentication. When Knox handles multiple concurrent requests, libpam might crash, causing Knox service to crash and requests to the Knox cdp-proxy-api topology to fail.

Use AD/LDAP authentication instead of PAM authentication for the cdp-proxy-api topology. In Cloudera Manager, go to the Knox service, select Configuration, and configure the gateway_api_authentication_provider parameter to use AD/LDAP authentication. For configuration instructions, see [Configure Apache Knox authentication for AD/LDAP](#).

Known Issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-76490: Ranger API bulk resource deletion fails when proxied through Knox

7.3.1 and its CHF, 7.3.2

Ranger API bulk resource deletion fails when the request is proxied through Knox. When Ranger sends a DELETE request with a body for bulk resource deletion, Knox does not forward the request body according to RFC 9110, causing the operation to fail.

None.

CDPD-76294: Knox service can not be started in a large size Private Cloud Base Cloudera Manager cluster

7.3.1.100 through 7.3.1.700, 7.3.2

For a large size Private Cloud Base Cloudera Manager cluster installed with Cloudera Runtime 7.3.1.0, you might face the problem that Knox can not be started with the following error message:

```
Wait Until Knox Gateway Can Serve Requests failed on Knox Gateway
```

Increase the Knox configuration parameter Knox Gateway Initial/Max Heapsize from 1 GiB to 2 GiB or 4 GiB, depending on the cluster size. Then save changes and run Restart Stale Services. After these steps, the Knox service can be started.

CDPD-71751: Creation of alias from the Cloudera Manager UI fails on FIPS

7.1.9 SP1 and its CHF, 7.3.1 and its CHF, 7.3.2

Users attempting to create aliases through the Cloudera Manager UI face issues in FIPS.

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

1. ssh to Knox host.
2. Export these directories: export KNOX_GATEWAY_DATA_DIR="/var/lib/knox/gateway/data"; export KNOX_GATEWAY_CONF_DIR="/var/lib/knox/gateway/conf"
3. Set the FIPS-specific options for the Knox CLI:

```
export KNOX_CLI_MEM_OPTS="--add-exports=java.base/sun.security.provider=bctls --add-exports=java.base/sun.security.provider=com.safelogic.cryptocomply.fips.core --add-modules=com.safelogic.cryptocomply.fips.core --add-modules=bctls --module-path=<BCTLS_JARS_DIR> -Dcom.safelogic.cryptocomply.fips.approved_only=true"
```

<BCTLS_JARS_DIR> is the directory containing the SafeLogic bctls and fips core jar files.

4. Run the following command to create the alias: /opt/cloudera/parcels/CDH/lib/knox/bin/knoxcli.sh create-alias <ALIAS_NAME> <ALIAS_VALUE>
5. Verify the addition using /opt/cloudera/parcels/CDH/lib/knox/bin/knoxcli.sh list-alias.

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

CDPD-71305: Concurrent impala shell connection failure

7.1.9 SP1 and its CHF, 7.3.1 and its CHF, 7.3.2

If a user makes a concurrent impala-shell connection through Knox, then the connection fails.

Use only one Knox role.

CDPD-3125: Logging out of Atlas does not manage the external authentication

7.1.9, 7.2.18, 7.3.1 and its CHF, 7.3.2

At this time, Atlas does not communicate a log-out event with the external authentication management, Apache Knox. When you log out of Atlas, you can still open the instance of Atlas from the same web browser without re-authentication.

To prevent additional access to Atlas, close all browser windows and exit the browser.

CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes

7.1.9, 7.3.1 and its CHF, 7.3.2

You must use a single Knox node to access Impala UI.

CDPD-22785: Improvements and issues need to be addressed in convert-topology kxcli command

7.1.9, 7.3.1 and its CHF, 7.3.2

None.

Knox issue with JDK version

7.1.9, 7.3.1 and its CHF, 7.3.2

jdk-1.8.0_391 is not supported.

Cloudera recommends using Cloudera supported JDKs.

CDPD-84236: Token generated by one Knox host fails with Unknown token error on another Knox host in Data Engineering High Availability clusters

7.3.1.400 through 7.3.1.700, 7.3.2

In Data Engineering High Availability clusters, a token generated by one Knox host may fail with an Unknown token error when accessed through another Knox host. This issue occurs due to a race condition in the PostgreSQL database, which prevents one of the Knox instances from properly initializing its configured token state service.

Restart Knox on all hosts.

Known Issues in Apache Kudu

Known issues and technical limitations for Kudu are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified Before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

You get "The user 'kudu' is not part of group 'hive' on the following hosts: " warning by the Host Inspector

If you are using fine grained authorization for Kudu, and you are also using Kudu-HMS integration with HDFS-Sentry sync, then you may get the "The user 'kudu' is not part of group 'hive' on the following hosts: " warning while upgrading.

Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

Fixed Issues in Livy

Review the list of known issues in Livy in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

There are no known issues carried forward that affect Livy in Cloudera Runtime 7.3.2.

Known Issues in Navigator Encrypt

Learn about the known issues in Navigator Encrypt, the impact or changes to the functionality, and the workaround.

Known Issues in Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

Failed to start navencrypt-mount.service

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 and its SPs and CHFs

After installing NavEncrypt on SLES, if the command "systemctl status navencrypt-mount" fails with the error: "Failed to start navencrypt-mount.service: Unit navencrypt-mount.service failed to load: No such file or directory", the systemd files need to be reloaded..

Run the command :

```
(sudo) systemctl daemon-reload
```

Known Issues in Oozie

Learn about the known issues in Oozie, the impact or changes to the functionality, and the workaround.

Known Issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known Issues identified before Cloudera Runtime 7.3.2

CDPD-93327: Passwords may be printed in the atlas-application.properties file

7.3.2

The /etc/atlas/conf/atlas-application.properties file is a publically readable file, and it may contain the Java trust store password. This file is read by the Oozie server and is propagated to the running actions. This allows the trust store password in the /etc/atlas/conf/atlas-application.properties file to be printed in the Oozie Launcher AM's (Yarn application) console logs, and the whole atlas-application.properties file is copied into the Oozie Launcher AM's (Yarn application) local directory.

None

Oozie jobs fail (gracefully) on secure YARN clusters when JobHistory server is down

7.3.2

If the JobHistory server is down on a YARN (MRv2) cluster, Oozie attempts to submit a job, by default, three times. If the job fails, Oozie automatically puts the workflow in a SUSPEND state.

Workaround: When the JobHistory server is running again, use the resume command to inform Oozie to continue the workflow from the point at which it left off.

Known Issues in Ozone

Known issues and technical limitations for Ozone are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2

CDPD-63350: Force deleting a FSO bucket and its contents while running `rb --force` from AWS S3 API is failing

7.3.2

Force deleting a FSO bucket and its contents while running the `rb --force` from the AWS S3 API might fail with an error, because the S3 client sends individual delete requests for each key in the bucket. It may delete or fail to delete individual keys or directories, depending on the availability of leaf elements. It can completely delete the bucket only when all the keys or directories are cleaned up, if some keys are not deleted, the bucket will not be deleted.

Sample error message:

```
# aws s3 rm s3://buck-fso --recursive
delete: s3://buck-fso/dir1/
delete: s3://buck-fso/dir1/dir2/
delete: s3://buck-fso/dir3/dir4/dir5/
delete: s3://buck-fso/dir3/dir4/
delete failed: s3://buck-fso/dir3/
# Rerun the same command again
# aws s3 rm s3://buck-fso --recursive
delete: s3://buck-fso/dir1/
delete: s3://buck-fso/dir3/
# To Confirm
# ozone sh key list s3v/buck-fso
[ ]
```

Run the `rb --force` command multiple times to completely cleanup the keys and directories.

Apache JIRA: [HDDS-9637](#)

CDPD-98751: Mismatched Replicas tab in the Recon UI fails to display containers with inconsistent replica checksums

7.3.2

In the Ozone Recon UI, Mismatched Replicas tab does not update or display containers when one or more replicas have differing checksums. Instead, they are displayed in the Under-Replicated tab.

The replica can be checked in the Under-Replicated tab, or can be cross-checked against API response.

CDPD-97512: Mismatch in Open Key count between Overview page and OM DB Insight in the Recon UI

7.3.2

In the Ozone Recon UI, the Open Key count on the Overview page (Summary section) may not match the count on the OM DB Insight Open Key tab. Users might see different Open Key values for the same cluster across these two views.

Use the OM DB Insight Open Key tab for the most accurate Open Key count.

CDPD-97376: Container replication counts mismatch between new and old Container pages in the Recon UI

7.3.2

In the Ozone Recon UI, Container replication counts—including Under-Replicated, Over-Replicated, and Mis-Replicated—differ between the new and old Container pages for the same cluster.

Refer to the old Container page to view accurate replication counts.

CDPD-97311: Incorrect Creation Time and Modification Time displayed on the Namespace Usage page in the Recon UI

7.3.2

In the Ozone Recon UI, the Namespace Usage page displays incorrect Created and Modified timestamps. These values do not accurately reflect the actual creation or last modification times or dates of the namespace, resulting in inaccurate metadata information.

Retrieve the correct timestamp values directly from the API response.

CDPD-97312: Mismatch between cluster State Container count and Container Summary totals

7.3.2

Ozone Recon does not sync all container states and can have discrepancies between the Storage Container Manager (SCM) container count and the Recon container count due to QUASI_CLOSED and other container states.

There is no workaround from the Ozone Recon. Use `ozone admin container report` CLI to get the correct container counts for various states in the Ozone cluster.

CDPD-99248: After cdh upgrade, Ozone encountered failure while executing the Finalize Upgrade for SCM on role Storage Container command

7.3.2

When finalizing an Ozone upgrade for the first time from Cloudera Manager, the finalize command `Finalize Upgrade for SCM on role Storage Container` may fail with the following message in stderr:

```
"Invalid response from Storage Container Manager.  
Current finalization status is: FINALIZATION_IN_PROGRESS"
```

Despite this, upgrade finalization is still running on the Storage Container Manager (SCM).

Even though the Cloudera Manager command failed, finalization is still running on SCM as the message indicates. Use the command `ozone admin scm finalizationstatus` to check the status of SCM finalization and wait for it to complete even if the Cloudera Manager command fails.

CDPD-98892: File Size Distribution bucket size range calculation is not correct

7.3.2

There are large number of buckets within a volume. The Ozone Recon UI might display incorrect file size distribution bucket size range calculation in the File Size Distribution chart in the Insights page.

None

Apache JIRA: [HDDS-14827](#)

CDPD-93116: Ozone client hangs for approximately five minutes intermittently when the disk is full

7.3.2

When the Ozone client writes data to the DataNode and there is an exception due to disk full condition or other error on the DataNode, then the client hangs for approximately five minutes due to continuous retry to write data till pipeline over the DataNode is closed.

Retry of the requests must be controlled using below configuration from the client side:

Table 7:

| Configuration | Recommended value |
|---|-------------------|
| hdds.ratis.raft.client.rpc.request.timeout | 30s |
| hdds.ratis.client.multilinear.random.retry.policy | 1s, 1 |
| hdds.ratis.client.exponential.backoff.max.sleep | 5s |
| hdds.ratis.client.exponential.backoff.base.sleep | 1s |
| hdds.ratis.client.exponential.backoff.max.retries | 2 |

Apache JIRA: [HDDS-14040](#)

Known Issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-91562: test_validate_certs_configs configuration is failing with the maximum lifetime validation

7.3.2, 7.3.1.600

In daylight saving time zones, the autogenerated Ozone certificate duration might differ from the expected duration. This discrepancy is minor, because the default certificate duration is 365 days or five years, depending on the Ozone component.

CDPD-75954: The ozone debug ldb command and ozone auditparser fails with java.lang.UnsatisfiedLinkError

7.3.2, 7.3.1.400

For information on workaround, see the [Changing temporary path for Ozone services and CLI tools](#) documentation.

CDPD-54885: Ozone Prometheus does not work with TLS

7.3.2, 7.3.1 and it's higher versions

The Prometheus service shipped by Ozone does not support TLS mode. So, Prometheus is not able to gather metrics from Ozone endpoints when TLS is enabled.

Go to Ozone Configuration Ozone Prometheus Endpoint Token and update the Ozone Prometheus Endpoint Token with any random string. This makes a token available in the process directory of the Ozone endpoint in plaintext, which Prometheus can use to get around the TLS limitation.

CDPD-56684: Keys get deleted when you do not have permission on volume

7.3.2, 7.3.1 and it's higher versions

When a volume is deleted, it recursively deletes the buckets and keys inside it and only then deletes the volume. The volume delete ACL check is done only in the end, due to which you may end up deleting all the data inside the volume without having delete permission on the volume.



Note: There was no bucket/key permission available which allowed the user to delete them recursively.

CDPD-50610: Large file uploads are slow with OPEN and stream data approach

7.3.2, 7.3.1 and it's higher versions

Hue file browser uses the append operation for large files. This API is not supported by Ozone in 7.1.9 and therefore large file uploads can be slow or timeout from the browser.

Use native Ozone client to upload large files instead of the Hue file browser.

OPSAPS-66469: Ozone-site.xml is missing if the host does not contain HDFS roles

7.3.2, 7.3.1 and it's higher versions

The client side ozone-site.xml (/etc/hadoop/conf/ozone-site.xml) is not generated by Cloudera Manager if the host does not have any HDFS role. Because of this, issuing Ozone commands from that host fails because it cannot find the service name to host name mapping. The error message is similar to this: # ozone sh volume list o3://ozoneabc 23/03/06 18:46:15 WARN ha.OMProxyInfo: OzoneManager address ozoneabc:9862 for serviceID null remains unresolved for node ID null Check your ozone-site.xml file to ensure ozone manager addresses are configured properly.

Add the HDFS gateway role on that host.

CDPD-74016: Running Ozone sh token print on token generated using ozone dtutil command results in Null Pointer Exception.

Print the token generated by dtutil using dtutil.

CDPD-74013: Ozone dtutil get token fails when using o3 or ofs schemas.

Use ozone sh token get to get a token for ozone file system.

CDPD-63144: Key rename inside the FSO bucket fails and displays the Failed to get parent dir error. This happens when running impala workloads with ozone.

None.

CDPD-74331: Key put fails and displays the Failed to write chunk error when there is a volume failure during configuration.

None.

CDPD-74475: YCSB test with Hbase on Ozone degrades performance.

7.3.2, 7.3.1 and it's higher versions

None.

CDPD-74884: Exclusive size of snapshot is always 0 when you run the info command on the Ozone snapshots. It is a statistics issue and does not impact the functionality of Ozone snapshot.

7.3.2, 7.3.1 and it's higher versions

The exclusiveSize and exclusiveReplicationSize stats presented by snapshot info are 0 even though the snapshot contains exclusive keys or files that are not present in other snapshots.

None.

Apache JIRA: [HDDS-11528](#)

CDPD-75042: AWS CLI rm or delete command fails to delete all files and directories on the Ozone FSO bucket. It only deletes the leaf node.

7.3.2, 7.3.1 and it's higher versions

None.

CDPD-75204: Namenode restart fails after dfs.namenode.fs-limits.max-component-length is set to a lower value and there is existing data present which exceeds the length limit.

7.3.2, 7.3.1 and it's higher versions

Increase the value for the dfs.namenode.fs-limits.max-component-length parameter and restart the namenode.

CDPD-75635: Ozone write fails intermittently as SCM remains in safemode.

7.3.2, 7.3.1 and it's higher versions

You must wait for SCM to come out of samemode or exit from safemode through CLI options.

Known Issues in Apache Parquet

Learn about the known issues in Parquet, the impact or changes to the functionality, and the workaround.

There are no new known issues identified in this release.

Known Issues in Apache Phoenix

Learn about the known issues in Phoenix, the impact or changes to the functionality, and the workaround.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

There are no known issues carried forward that affect Phoenix in Cloudera Runtime 7.3.2.

Known Issues in Apache Ranger

Learn about the known issues in Apache Ranger, the impact or changes to the functionality, and the workaround.

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known Issues identified before Cloudera Runtime 7.3.2

OPSAPS-75673: Wrong enablement of Ranger RMS Database Full Sync command

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs, 7.1.9 and its SPs and CHFs, 7.1.8 and its SPs and CHFs

The `Ranger RMS Database Full Sync` command should be enabled only when all RMS server instances are stopped. This is required to ensure that the RMS database synchronizes correctly without introducing conflicts or data corruption. However, when HA (High Availability) is enabled on the cluster, the command becomes available from Cloudera Manager Ranger RMS Actions drop-down, even though only one Ranger RMS instance is stopped while the others are still running.

None.

CDPD-68806: The Revoke operation for users belonging to a group or role permission does not function as expected

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

List command is listing all the tables even when the user permission is revoked. And also the command does not add any deny policy to Ranger for that specific user.

This behavior is currently not supported in HBase shell. Must be handled manually using the Ranger policy change.

CDPD-68739: The revoke command does not work when using the HBase shell

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

While using the HBase shell, running the revoke command does not cancel the user permission. Users are able to perform actions even after running the revoke command.

None.

CDPD-58704: hadoop roll key/key delete command shows operation failed error when one KMS host is down, even when operation succeeds

7.3.2, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its CHF, 7.1.9 and its SPs and CHF

In case of rollover/delete, client sends one more (last after delete request) request to KMS instances to clean their cache and that too to all registered kms instances. if one KMS instance is stopped (not deleted), the client gets a runtime exception.

This simply returns the runtime exception on client end for stopped instances but doesn't break any functionality.

CDPD-56803: When there is no existing policy for user and a revoke request comes from hbase, then will get this error

7.3.2, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its CHF, 7.1.9 and its SPs and CHF

None.

CDPD-56741: Improvement in log message when jwtauth not used

7.3.2, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its CHF, 7.1.9 and its SPs and CHF

None.

CDPD-56738: Ranger RMS showing FileNotFoundException: /usr/share/java/oraclepki.jar in Oracle 19 setup

7.3.2, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its CHF, 7.1.9 and its SPs and CHF

This is a warning log printed in catalina.out file when Ranger RMS server is initialized. The following exception is observed only in Oracle 19 setup: FileNotFoundException: /usr/share/java/oraclepki.jar

None.

CDPD-48975: Ranger KMS KTS to KMS DB migration: keys with the same name but different case are not migrated

7.3.2, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its CHF, 7.1.9 and its SPs and CHF

KMS keys are not case sensitive.

No workaround. Such key combinations are very rare and the migration doc was updated to check such keys before starting the migration.

CDPD-42598: Kafka policy creation allowed with incorrect permissions

7.3.2, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its CHF, 7.1.9 and its SPs and CHF, 7.1.8 and its SPs and CHF

When creating a Kafka policy from the UI, the permissions "Idempotent write" and "Cluster action" are not displayed as they are not applicable for the "topic" resource, but when creating a policy for the "topic" resource with the permissions "Idempotent write" and "Cluster Action", the policy is created successfully when the expected behaviour is that the policy creation must fail as the permission is not applicable for the Kafka topic resource.

None.

CDPD-41582: Atlas Resource Lookup : Classification for "entity-type" lists only classification for the following payload: {"resourceName": "classification", "userInput": "", "resources": {"classification": []}}

7.3.2, 7.3.1 and its SPs and CHF, 7.1.9 SP1 and its CHF, 7.1.9 and its SPs and CHF, 7.1.8 and its SPs and CHF

Expectation is to return all the classifications. But the response has only "classification". Happens similarly for entity-label, entity-business-metadata.

None.

Known Issues in Ranger KMS

Learn about the known issues in Ranger KMS, the impact or changes to the functionality, and the workaround.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

There are no known issues carried forward that affect Ranger KMS in Cloudera Runtime 7.3.2.

Known Issues in Schema Registry

Known issues and technical limitations for Schema Registry are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-40380: Authorization checking issue when Kerberos is disabled

7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

Due to an issue in Ranger, when Kerberos is disabled then it is not possible to check authorization.

1. Open Schema Registry configuration in Cloudera Manager.
2. Find the `ranger.plugin.schema-registry.service.name` field.
3. Replace `GENERATED_RANGER_SERVICE_NAME` with the actual name of the service.
4. Restart the Schema Registry service.

CDPD-49304: AvroConverter does not support composite default values

7.1.7 and all its SP and CHF releases, 7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

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

7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

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

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

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

7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

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

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

Known Issues in Apache Solr

Learn about the known issues in Apache Solr, the impact or changes to the functionality, and the workaround.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

There are no known issues carried forward that affect Solr in Cloudera Runtime 7.3.2.

Known Issues in Spark

Review the list of known issues in Spark in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2

CDPD-67806: Migration of ORC table with timestamp column to Iceberg with `spark.sql.timestampType=TIMESTAMP_NTZ` fails

7.2.18, 7.3.1.x

The Hive External Catalog converts NTZ timestamp types back to `TimeStampType`, which Hive does not support, leading to migration failures.

Set `spark.sql.iceberg.use-timestamp-without-timezone-in-new-tables` and `spark.sql.iceberg.handle-timestamp-without-timezone` to `true` to migrate ORC tables to Iceberg with timestamp columns without timezone..

CDPD-90447: Table is stored in HDFS when cluster is enabled with RAZ and S3 is set to default storage system

7.3.2

Tables created on a RAZ-enabled clusters with S3 set as the default storage system are stored in HDFS instead of the specified S3 bucket.

None. S3 was removed as the default storage.

OPSAPS-75684: Spark fails due to Zookeeper Custom Kerberos Principal issue

7.3.2

Incorrect Zookeeper principal configuration and missing JVM property setup leads to SASL authentication failures.

When a custom Zookeeper principal is used, add the `-Dzookeeper.sasl.client.username=[***USE RNAME***]` JVM argument to `spark.*.defaultJavaOptions` or `spark.*.extraJavaOptions` in `spark-defaults.conf`.

Known issues identified before Cloudera Runtime 7.3.2

CDPD-95322: Missing Atlas lineage for Spark Iceberg tables from MERGE INTO

7.3.2, 7.3.1.0 and higher

Spark SQL MERGE INTO statements on Iceberg tables are not transmitting lineage data to Atlas.

None.

CDPD-94393: RuntimeWarning Failed to add file" message appears even when Spark successfully loads files

7.3.2, 7.3.1.0 and higher

In both Spark 2 and 3, due to an exception when attempting to add files to the Python path, the RuntimeWarning: Failed to add file message appears even when the Python JAR file is successfully loaded.

None. You can safely ignore the message as the file is loaded successfully and the message does not affect job completion.

Spark 3: RAPIDS Accelerator is not available

7.3.2, 7.3.1.0 and higher

The RAPIDS Accelerator for Apache Spark is currently not available in Cloudera Runtime 7.3.1

None.

The CHAR(n) type handled inconsistently, depending on whether the table is partitioned or not.

7.3.1

7.3.2, 7.3.1.100 CHF1 and higher

In upstream Spark 3 the spark.sql.legacy.charVarcharAsString configuration was introduced, but it does not solve all incompatibilities with Spark 2.

None. A new configuration spark.cloudera.legacy.charVarcharLegacyPadding will be introduced in a future version to keep compatibility with Spark 2, but it isn't available in 7.3.1.



Note: The CHAR type is legacy in SQL, and using it is discouraged. Cloudera recommends using VARCHAR or STRING instead.

Apache Jira: [SPARK-33480](#)

Known Issues in Spark Atlas Connector

Review the list of known issues in Spark Atlas Connector in Cloudera Runtime 7.3.2.0, its service packs and cumulative hotfixes.

Known Issues Identified In Cloudera Runtime 7.3.2.0

There are no known issues identified in this release.

Known Issues Identified Before Cloudera Runtime 7.3.2.0

There are no new known issues identified in 7.3.1.x.

Known Issues for Sqoop

Known issues and technical limitations for Sqoop are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified Before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-54770: Unable to read Sqoop metastore created by an older HSQLDB version

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

If you have upgraded to Cloudera Base on premises 7.1.8 Cumulative hotfix 4 or higher versions, you may encounter issues in reading the Sqoop metastore that was created using an older version of HyperSQL Database (HSQLDB).

Cloudera upgraded the HSQLDB dependency from 1.8.0.10 to 2.7.1 and this causes incompatibility issues in Sqoop jobs that are stored in HSQLDB.

After upgrading to Cloudera Base on premises 7.1.8 Cumulative hotfix 4, you must upgrade the Sqoop metastore and convert the database files to a format that can easily be read by HSQLDB 2.7.1. For more information, see [Troubleshooting Apache Sqoop issues](#).

CDPD-44431: Using direct mode causes problems

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

Using direct mode has several drawbacks:

- Imports can cause an intermittent and overlapping input split.
- Imports can generate duplicate data.
- Many problems, such as intermittent failures, can occur.
- Additional configuration is required.

Stop using direct mode. Do not use the `--direct` option in Sqoop import or export commands.

Sqoop direct mode is disabled by default. However, if you still want to use it, enable it by either setting the `sqoop.enable.deprecated.direct` property globally in Cloudera Manager for Sqoop or by specifying it in the command-line through `-Dsqoop.enable.deprecated.direct=true`.

CDPD-3089: Avro, S3, and HCat do not work together properly

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs, 7.1.7 SP1 and its CHFs

Importing an Avro file into S3 with HCat fails with Delegation Token not available.

Parquet columns inadvertently renamed

Problem: Column names that start with a number are renamed when you use the `--as-parquetfile` option to import data.

Prepend column names in Parquet tables with one or more letters or underscore characters.

Importing Parquet files might cause out-of-memory (OOM) errors

Problem: Importing multiple megabytes per row before initial-page-run check (ColumnWriter) can cause OOM. Also, rows that vary significantly by size so that the next-page-size check is based on small rows, and is set very high, followed by many large rows can also cause OOM.

PARQUET-99

Known issues in Streams Messaging Manager

Known issues and technical limitations for Streams Messaging Manager are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known Issues identified in Cloudera Runtime 7.3.2

Alert policies based on the LEADER ELECTION PER SEC attribute do not trigger

7.3.2

If the monitored Kafka cluster is running in KRaft mode, alert policies based on the **LEADER ELECTION PER SEC** alert policy attribute will not trigger because the underlying metrics are not available from Kafka. The attribute remains visible and selectable in Streams Messaging Manager even if the monitored cluster is running in KRaft mode.

None.

CDPD-99371: Prometheus metrics do not function properly with KRaft mode

7.3.2

If the Kafka cluster monitored by Streams Messaging Manager is running in KRaft mode and Streams Messaging Manager is configured to use Prometheus as its metrics store, alerts related to the active controller do not work as expected. In KRaft mode, the `broker_activecontrollercount` metric is reported by KRaft controllers rather than brokers. This change in metric reporting can cause aggregation issues and lead to alerts triggering indefinitely. Additionally, adding the `/api/prometheus-metrics` endpoint from KRaft hosts to the Prometheus configuration can cause anomalies in other alerts.

Use Cloudera Manager as the metrics store for KRaft-based Kafka clusters. Alternatively, if using Prometheus, disable all alerts based on the `broker_activecontrollercount` metric.

Known issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

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

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

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

View the Streams Messaging Manager UI logs on the host.

CDPD-39313: Some numbers are not rendered properly in Streams Messaging Manager UI

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

Very large numbers can be imprecisely represented on the UI. For example, bytes larger than 8 petabytes would lose precision.

None.

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

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

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

Streams Messaging Manager cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You need to override bootstrap server URL (hostname:port as set in the listeners for broker). Add the bootstrap server details in Streams Messaging Manager safety valve in the following path:

1. In Cloudera Manager, go to `SMMConfigurationStreams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml`.

2. Add the following value for bootstrap servers.

```
streams.messaging.manager.kafka.bootstrap.servers=<comma-separated list of brokers>
```

3. Save your changes.
4. Restart Streams Messaging Manager.

CDPD-82560: Streams Messaging Manager known unprotected endpoint

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

The Streams Messaging Manager Frontend exposes the /cm-configs and /configs public endpoints without requiring authentication. These endpoints do not share any sensitive information.

None.

Known Issues in Streams Replication Manager

Known issues and technical limitations for Streams Replication Manager are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-22089: Streams Replication Manager does not sync re-created source topics until the offsets have caught up with target topic

7.1.7 and all its SP and CHF releases, 7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

Messages written to topics that were deleted and re-created are not replicated until the source topic reaches the same offset as the target topic. For example, if at the time of deletion and re-creation there are a 100 messages on the source and target clusters, new messages will only get replicated once the re-created source topic has 100 messages. This leads to messages being lost.

None

CDPD-11079: Blacklisted topics appear in the list of replicated topics

7.1.7 and all its SP and CHF releases, 7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

If a topic was originally replicated but was later disallowed (blacklisted), it will still appear as a replicated topic under the /remote-topics REST API endpoint. As a result, if a call is made to this endpoint, the disallowed topic will be included in the response. Additionally, the disallowed topic will also be visible in the Streams Messaging Manager UI. However, its Partitions and Consumer Groups will be 0, its Throughput, Replication Latency and Checkpoint Latency will show N/A.

None

CDPD-30275: Streams Replication Manager may automatically re-create deleted topics on target clusters

7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

If auto.create.topics.enable is enabled, deleted topics might get automatically re-created on target clusters. This is a timing issue. It only occurs if remote topics are deleted while the replication of the topic is still ongoing.

1. Remove the topic from the topic allowlist with srm-control. For example:

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGET_CLUSTER] --remove [TOPIC1]
```

2. Wait until Streams Replication Manager is no longer replicating the topic.
3. Delete the remote topic in the target cluster.

CDPD-80872: Streams Replication Manager replication-records-lag is incorrect on empty partitions with non-zero end offset

7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

When a replicated partition is empty and its end offset is non-zero, the replication-records-lag metric is incorrectly reported as the end offset (instead of 0 or NaN).

Wait for a new message to be written into the partition. When it gets replicated, the metric is reported correctly.

Limitations

Streams Replication Manager cannot replicate Ranger authorization policies to or from Kafka clusters

Due to a limitation in the Kafka-Ranger plugin, Streams Replication Manager cannot replicate Ranger policies to or from clusters that are configured to use Ranger for authorization. If you are using Streams Replication Manager to replicate data to or from a cluster that uses Ranger, disable authorization policy synchronization in Streams Replication Manager. This can be achieved by clearing the Sync Topic Acls Enabled (sync.topic.acls.enabled) checkbox.

Known Issues in YARN, YARN Queue Manager and MapReduce

Known issues and technical limitations for YARN and YARN Queue Manager are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-75652: Reverse DNS lookup fails for YARN but works for HDFS

7.3.2, 7.3.1 and its SPs and CHFs

When submitting a YARN application from a host without the correct DNS configuration (reverse DNS does not work for the YARN ResourceManager's host) Kerberos principal validation error, such as Server has invalid Kerberos principal. This is because the reverse DNS configuration does not work for YARN ResourceManager's host.

Add the following to YARN Service Advanced Configuration Snippet (Safety Valve) for yarn-site.xml:

```
<property>
<name>yarn.resourcemanager.principal.pattern</name>
<value>*</value>
</property>
```

COMPX-14682: Fix health check after Queue Manager restart

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

In some cases, after the QM webapp and config-service are restarted Cloudera Manager reports a healthy status after some time. However, the configuration service might not be ready and usable yet when this status is reported. Therefore, QM (webapp) is not ready and usable yet.

Wait for a minute or two after restart and then try QM.

COMPX-14820: Delete Queue and its Children throws "Queue capacity was reduced to zero, but failed to delete queue."

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

When trying to perform the operation "Delete Queue and its Children" on a queue that has one or more siblings, the operation fails as YARN has some constraints. If the queue performing the operation "Delete Queue and its Children" is a leaf node, then the operations succeeds.

None.

A fresh install of 7.3.1 or its SPs or CHFs does not allow user to bypass the Setup Database screen for YARN Queue Manager

7.3.2, 7.3.1 and its SPs and CHFs

YARN Queue Manager in Cloudera Base on premises 7.3.1 does not require you to install a PostGres database, therefore users should not see the Setup Database screen and should be able to skip the Setup Database screen. With this known issue, users who are conducting a fresh install of 7.3.1 or its SPs or CHFs are not able to bypass the Setup Database screen as expected.

1. When conducting a fresh install of YARN Queue Manager in 7.3.1 or its SPs or CHFs, you must ensure that you have both Cloudera and Cloudera Manager upgraded to 7.3.1.
2. When you reach the Setup Database screen in the Cloudera Manager installation wizard for Queue Manager, enter any dummy values for the following fields:

- a. Database name: configstore
- b. Database Username: dbuser
- c. Database Password: dbpassword
- d. Database Hostname: localhost

YARN Queue Manager will not connect to PostGres with the above details and will fall back to the embedded database.

3. Run the following script command in a browser console to enable the Continue button:

```
document.querySelector('.btn.next').removeAttribute('disabled');
```

4. Click Continue and proceed with the YARN Queue Manager installation.
5. After installation is complete, SSH into the host that has Queue Manager installed, and run this command: `sed -i 's/migrationCompleted=true/migrationCompleted=false/' /var/lib/hadoop-yarn/migration.properties`



Note: Enable Queue Manager in the YARN configurations, and restart YARN.

6. Restart YARN Queue Manager.

20202 Database migration after enabling opt-in migration

7.3.2, 7.3.1 and its SPs and CHFs

When migrating from an H2 database to a PostgreSQL database in YARN Queue Manager after installation or upgrade, you might encounter an issue only when you have followed the following specific scenario:

- New install or upgrade to Cloudera 7.1.9, forcing migration from H2 to PostgreSQL database.
- Upgrade to Cloudera 7.1.9 CHF2, moving back to H2 database.
- Upgrade to Cloudera 7.1.9 SP1 with valid PostgreSQL connection details in Queue Manager configurations.

To avoid any issues during the upgrade to version Cloudera 7.1.9 SP1, ensure that PostgreSQL connection details are removed from the YARN database configuration if you prefer to continue using the H2 database.

Queue Manager does not open on using a custom user with a default Kerberos principal

7.3.2, 7.3.1 and its SPs and CHFs

If a custom user is used with the default Kerberos principal, the Queue Manager web UI displays an HTTP ERROR 400 error.

Ensure that the Queue Manager process_username property matches the YARN process_username property.

Third-party applications do not launch if MapReduce framework path is not included in the client configuration

7.3.2, 7.3.1 and its SPs and CHFs

MapReduce application framework is loaded from HDFS instead of being present on the NodeManagers. By default the mapreduce.application.framework.path property is set to the appropriate value, but third-party applications with their own configurations does not launch.

Set the mapreduce.application.framework.path property to the appropriate configuration for third-party applications.

JobHistory URL mismatch after server relocation

7.3.2, 7.3.1 and its SPs and CHFs

After moving the JobHistory Server to a new host, the URLs listed for the JobHistory Server on the ResourceManager web UI still point to the old JobHistory Server. This affects existing jobs only. New jobs started after the move are not affected.

For any existing jobs that have the incorrect JobHistory Server URL, there is no option other than to allow the jobs to roll off the history over time. For new jobs, make sure that all clients have the updated the mapred-site.xml file that references the correct JobHistory Server.

YARN cannot start if Kerberos principal name is changed

7.3.2, 7.3.1 and its SPs and CHFs

If the Kerberos principal name is changed in Cloudera Manager after launch, YARN does not start. In such cases, the keytabs can be correctly generated but YARN cannot access ZooKeeper with the new Kerberos principal name and old ACLs.

There are two possible workarounds:

- Delete the znode and restart the YARN service.
- Use the reset ZK ACLs command. This also sets the znodes below /rmstore/ZKRMStateRoot to world:anyone:cdw which is less secure.

YARN Resource Manager UI does not display logs

7.3.2, 7.3.1 and its SPs and CHFs

When the parameter mapreduce.cluster.acl.enabled is set to true, the Yarn RM UI does display and logs and the Logs are not available message is displayed.

Set mapreduce.cluster.acl.enabled to false.

Known Issues in Apache ZooKeeper

Learn about the known issues in ZooKeeper, the impact or changes to the functionality, and the workaround.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

There are no new known issues from previous releases that affect ZooKeeper in Cloudera Runtime 7.3.2.

Behavioral Changes in Cloudera Runtime 7.3.2

Behavioral changes denote a marked change in behavior from the previously released version to this version of Cloudera Runtime.

Behavioral Changes in Atlas

Functional adjustments and behavioral updates for Atlas are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Atlas, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see [Behavioral Changes](#).

Cloudera Runtime 7.3.2

Summary: Automatic purging of soft-deleted entities is introduced

Previous behavior:

Previously, only the API call `DELETE /api/atlas/admin/purge` was available to manually purge soft-deleted entities. Additionally, the `DELETE api/atlas/v2/entity/guid/{guid}` API call could not delete the column lineages entities of Hive, Impala and Spark process entities. This could lead to sparse graphs resulting in reduced query performance.

New behavior:

A built-in auto-purge mechanism is introduced, deleted entities are purged in two stages. The first stage is a soft-delete at each Atlas startup. In the second-stage, soft-deleted process entities are purged based on a cron job. For more information, see [Atlas Auto-Purging overview](#).

Summary:

Entity attributes details and sparkPlanDescription are no longer sent in the Spark process entity

Previous behavior:

The spark_process entity attributes details and sparkPlanDescription are populated with query plan details, which can contain a large amount of text, often in megabytes. This amount of data can incur unnecessary processing costs.

New behavior:

The atlas.spark.plan.enabled is set to false by default. Set it to true to send the details and sparkPlanDescription attributes in the Spark process entity. When these attributes are not sent, the cost of having large amount of data processed in Atlas is avoided.

Behavioral Changes in Avro

Functional adjustments and behavioral updates for Avro are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

There are no behavioral changes in this release.

Behavioral Changes in Cloud Connectors

Behavioral changes denote a marked change in behavior from the previously released version to this version of Cloud Connectors.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Cloud Connectors, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.700. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see [Behavioral Changes](#).

There are no behavioral changes in this release.

Behavioral Changes in Cruise Control

Functional adjustments and behavioral updates for Cruise Control are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

Summary: ZooKeeper service dependency removed from Cruise Control

Previous behavior:

Cruise Control had a ZooKeeper service dependency in Cloudera Manager.

New behavior:

The ZooKeeper service dependency is removed.

Summary: Default JMX settings changed to restrict connections to localhost

Previous behavior:

No dedicated Cloudera Manager configuration option existed to control JMX JVM flags for Cruise Control. JMX connections were unrestricted by default.

New behavior:

A new `CC_JMX_OPTS` Cloudera Manager configuration option is available. Its default value restricts JMX connections to localhost only and enables SSL. If you previously set `CC_JMX_OPTS`, your custom value is preserved on upgrade. To revert to open JMX, update `CC_JMX_OPTS` in Cloudera Manager.

Summary: Default Supported Goals and Anomaly Detection Goals updated in Cloudera Manager

Previous behavior:

The Cloudera Manager defaults for Supported Goals did not include `BrokerSetAwareGoal`. The defaults for Anomaly Detection Goals covered rack, replica, and disk capacity only, not network inbound, network outbound, or CPU capacity goals.

New behavior:

The default Supported Goals list includes `BrokerSetAwareGoal`. The default Anomaly Detection Goals list also includes `NetworkInboundCapacityGoal`, `NetworkOutboundCapacityGoal`, and `CpuCapacityGoal`. Cloudera Manager defaults for these goal lists now match the Cruise Control defaults. No action is required after upgrade.

Behavioral Changes in HBase

Functional adjustments and behavioral updates for HBase are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for HBase, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.700. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see [Behavioral Changes](#).

Cloudera Runtime 7.3.2

Summary: HBOSS related configuration items are removed.

Previous behavior:

HBase contains the following HBOSS related configuration items.

- fs.hboss.fs.s3a.impl
- fs.hboss.sync.impl

New behavior:

The above configuration items are removed because HBOSS support is deprecated.

Summary: Removed the hbase.secure.rpc.engine configuration property

Previous behavior:

The hbase.secure.rpc.engine configuration property is obsolete because HBase performs replication securely by default.

New behavior:

The hbase.secure.rpc.engine configuration property is removed because it is obsolete and no longer needed.

Summary: The default values for the following configuration items are updated

Previous behavior:

| Parameter name | Description | Default value |
|---|--|---------------|
| hbase.bucketcache.minfactor | The minimum percentage of BucketCache usage that should be targeted by the mass eviction process. | - |
| hbase.server.netty.tls.client.auth.mode | The client authentication mode for mTLS authentication when TLS RPC is in use. Three modes are NEED, WANT and NONE. | NEED |
| ip_version | Specifies the IP version the service must use for network communication. <ul style="list-style-type: none"> • IPv4 - Uses IPv4 exclusively. • IPv6 - Uses IPv6 exclusively. • Dual-stack (IPv4 & IPv6) - Supports both IPv4 and IPv6, enabling communication over both protocols. | - |
| hbase.client.netty.tls.enabled | Encrypt communication between clients and HBase RPC client mode using Transport Layer Security (TLS) (formerly known as Secure Socket Layer (SSL)). | - |
| hbase.bucketcache.acceptfactor | The percentage of BucketCache usage that must trigger the eviction of blocks. | - |

| Parameter name | Description | Default value |
|-----------------------------------|--|---|
| hbase.rpc.tls.truststore.location | The location on disk of the trust store, in .jks format, used to confirm the authenticity of TLS/SSL servers that HBase Client RPC might connect to. This trust store must contain the certificate(s) used to sign the service(s) connected to. If this parameter is not provided, the default list of well-known certificate authorities is used instead. | CM_AUTO_TLS |
| fs.s3a.impl | The implementation to use for S3A FileSystem access. | org.apache.hadoop.hbase.oss.HBaseObjectStoreSemantics |
| hbase.rpc.tls.truststore.password | The password for the HBase Client RPC TLS/SSL Trust Store File. This password is not required to access the trust store; this field can be left blank. This password provides optional integrity checking of the file. The contents of trust stores are certificates, and certificates are public information. | CM_AUTO_TLS |

New behavior:

| Parameter name | Description | Default value |
|---|--|---------------|
| hbase.bucketcache.minfactor | The minimum percentage of BucketCache usage that should be targeted by the mass eviction process. | 0.95 |
| hbase.server.netty.tls.client.auth.mode | The client authentication mode for mTLS authentication when TLS RPC is in use. Three modes are NEED, WANT and NONE. | NONE |
| ip_version | Specifies the IP version the service must use for network communication. <ul style="list-style-type: none"> IPv4 - Uses IPv4 exclusively. IPv6 - Uses IPv6 exclusively. Dual-stack (IPv4 & IPv6) - Supports both IPv4 and IPv6, enabling communication over both protocols. | IPV4 |
| hbase.client.netty.tls.enabled | Encrypt communication between clients and HBase RPC client mode using Transport Layer Security (TLS) (formerly known as Secure Socket Layer (SSL)). | FALSE |
| hbase.bucketcache.acceptfactor | The percentage of BucketCache usage that must trigger the eviction of blocks. | 0.98 |
| hbase.rpc.tls.truststore.location | The location on disk of the trust store, in .jks format, used to confirm the authenticity of TLS/SSL servers that HBase Client RPC might connect to. This trust store must contain the certificate(s) used to sign the service(s) connected to. If this parameter is not provided, the default list of well-known certificate authorities is used instead. | None |

| Parameter name | Description | Default value |
|-----------------------------------|--|--|
| fs.s3a.impl | The implementation to use for S3A FileSystem access. | org.apache.hadoop.fs.s3a.S3AFileSystem |
| hbase.rpc.tls.truststore.password | The password for the HBase Client RPC TLS/SSL Trust Store File. This password is not required to access the trust store; this field can be left blank. This password provides optional integrity checking of the file. The contents of trust stores are certificates, and certificates are public information. | None |

Summary: Add security headers to Thrift/HTTP server

Previous behavior:

Previously, security headers were absent in the responses of Thrift or HTTP servers when utilizing HBase-exposed services.

New behavior:

To address this, additional security headers are introduced in the responses of the HBase Thrift server when HTTP or HTTPS transport is enabled.

Apache JIRA: [HBASE-27118](#)

Behavioral Changes in HDFS

Behavioral changes denote a marked change in behavior from the previously released version to this version of HDFS.

Cloudera Runtime 7.3.2

There are no behavioral changes in this release.

Behavioral Changes in Hive

Functional adjustments and behavioral updates for Hive are introduced in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Hive, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see [Behavioral Changes](#).

Summary:

Handling invalid date formats in to_date function.

The behavior of the to_date function has changed between Cloudera Runtime versions 7.1.7 SP2 and 7.1.9 when handling invalid date formats.

Previous behavior:

Valid dates (e.g., YYYY-MM-DD): Returned correct results.

Invalid dates: Returned random, unexpected dates instead of NULL.

New behavior:

Following the changes introduced in [HIVE-28483](#):

Valid dates: Continue to return correct results.

Invalid dates (e.g., DD-MM-YYYY): Now return NULL.

Summary:

Metastore secondary connection pool size is now configurable

Previous behavior: The Metastore's secondary connection pool had a fixed size of 2. This often led to connection limitations, especially under heavy workloads.

New behavior: You can now configure the metastore's secondary connection pool size using the property `datanucleus.connectionPool.secondary.maxPoolSize`. This lets you adjust the pool beyond its default of 2, preventing connection limitations and improving performance.

Summary:

New configuration available to disable the Partition Management Task

Previous behavior:

It was not possible to disable the `PartitionManagementTask`, which put a heavy load on the Cloudera Manager Metastore when managing a large number of tables and partitions. Customers had a tedious workaround by providing a pattern via `metastore.partition.management.database.pattern / metastore.partition.management.table.pattern`.

New behavior:

You can now set the `metastore.partition.management.task.frequency` configuration to 0 to disable the partition management task cluster-wide. This helps to reduce the load on the Cloudera Manager Metastore.

Apache Jira: [HIVE-25324](#)

Summary:

New default restriction for Iceberg table data file locations

Previous behavior:

In earlier versions, the `hive.iceberg.allow.datafiles.in.table.location.only` property was set to false by default. This allowed Hive to access and read Iceberg data files even if they were located outside of the specific table directory.

New behavior:

The default value for `hive.iceberg.allow.datafiles.in.table.location.only` is now true. This security enhancement ensures that only data files located within the table directory are accessible. If you attempt to read an Iceberg table that contains data files outside of its directory, Hive now returns an error. If your existing workflows rely on data files stored in external locations, you can disable this restriction by using the Cloudera Manager to set the property to false.

Summary:

Lineage information computation enabled by default

Previous behavior:

Previously, lineage information was computed only if specific hardcoded post-execution hooks were configured or if the deprecated `HIVE_LINEAGE_INFO` property was set to true.

New behavior:

Lineage information is now collected by default for all queries. However, the system only records and passes lineage to hooks for the specific query types defined in the `HIVE_LINEAGE_STATEMENT_FILTER` property. By default, this includes `CREATE_TABLE`, `CREATE_TABLE_AS_SELECT`,

`CREATE_VIEW`, `CREATE_MATERIALIZED_VIEW`, and `LOAD`.

Apache Jira: [HIVE-28768](#)

Summary:

Increased Batch Sizes for COMPUTE STATS

Previous behavior:

The `COMPUTE STATS` query previously failed on tables containing more than 5000 columns. This issue was specific to wide tables and could not be resolved by dropping and rerunning the query.

New behavior:

To resolve this, we enable the batch retrieval or insertion of the object metadata by default. The default value of the `hive.metastore.direct.sql.batch.size` property is changed from 0 to 1000, and the default value of the `metastore.rawstore.batch.size` property is changed from -1 to 500. After this change, `COMPUTE STATS` queries now run successfully on tables with more than 5000 columns.

Behavioral Changes in Cloudera Data Explorer (Hue)

Functional adjustments and behavioral updates for Cloudera Data Explorer (Hue) are introduced in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Cloudera Data Explorer (Hue), and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see [Behavioral Changes](#).

There are no behavioral changes introduced in this release.

Behavioral Changes in Impala

Functional adjustments and behavioral updates for Impala are introduced in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Impala, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see [Behavioral Changes](#).

Summary:

Impala no longer allows unauthenticated user impersonation through Knox JDBC gateway

Previous behavior:

When connecting to Impala using the Knox JDBC gateway URL (e.g., `jdbc:impala://[knoxhost]:[knoxport]/...`), you could specify any Impala user through the `impala.doas.user` property and successfully impersonate them, regardless of whether your Knox user had explicit Impala proxy user permissions. This unintended behavior bypassed security checks and allowed for delegation to users who should not have been accessible.

New behavior:

Knox JDBC Gateway authenticates as itself and delegates to the logged in user, so it has already performed some delegation. If `impala.doas.user` is passed by the client through Knox JDBC Gateway, it will now result in an error as we don't support secondary delegation.

Summary:

Query cancellation during frontend planning

Previous behavior:

Previously, you could not cancel queries while they were in the analysis or planning stages. If a query was waiting for metadata to be loaded from the Catalog Server or the Hive Metastore (HMS), it would continue to run until the frontend planning process was complete. You had to wait for the query to reach the execution stage before a cancellation request could take effect.

New behavior:

You can now cancel queries during the frontend planning and metadata operation stages. If a query is interrupted by a user or a timeout while in the planning stage, the system triggers an interruption to the specific thread. This is particularly effective when planning is waiting on external metadata services. The cancellation process blocks until the frontend reaches an interruption point and returns to the backend to finalize the query status.

Summary:

Default value for `fe_service_threads` increased to improve concurrency

Previous behavior:

The default value for the `fe_service_threads` setting was 64.

New behavior:

Starting with Cloudera on premises 7.3.2, the default value is 128.

Summary:

Cleanup subdirectories in truncate/insert overwrite if recursing listing is enabled

Previous behavior:

Impala did not consistently delete files located in subdirectories of external tables during `TRUNCATE` and `INSERT OVERWRITE` operations, even when recursive listing was enabled. This led to leftover data in subdirectories after these operations, resulting in data corruption.

New behavior:

After this change, directories are also deleted in addition to (non-hidden) data files, with the exception of hidden and ignored directories. Now, setting `DELETE_STATS_IN_TRUNCATE=false` is no longer supported by default when truncating non-transactional tables; attempting this will result in an exception. If the old behavior is absolutely required, you can set the `--truncate_external_tables_with_hms` flag to false, but be aware that this will also reintroduce the bug that was fixed by this change.

Apache Jira: [IMPALA-14189](#), [IMPALA-14224](#)

Summary:

Parquet late materialization behavior has changed

Previous behavior:

Parquet late materialization feature was disabled by default. You would use the `parquet_late_materialization_threshold` query option to set the minimum number of consecutive filtered rows required to trigger late materialization. The default value was -1. The feature was not supported for collection columns.

New behavior:

Parquet late materialization feature is enabled by default for all types including collections. The `parquet_late_materialization_threshold` is now set to 1 if the query option is greater than or equal to 0 and there is a collection value that can be skipped. Otherwise, the value is the same as the query option, which defaults to 20.

Apache Jira: [IMPALA-3841](#)

Summary:

TCP Keepalive is now enabled by default for client connections

Previous behavior:

TCP keepalive was disabled by default for client connections. Idle connections dropped by load balancers remained active in Impala, consuming service threads (fe_service_threads).

New behavior:

TCP keepalive is now enabled by default for all client connections, enhancing stability and availability. Impala is configured to check idle connections aggressively, every 10 minutes.

Apache Jira: [IMPALA-14031](#)

Summary:

Support for load-based routing in impala-proxy

Previous behavior:

The impala-proxy used a random selection policy to choose a coordinator. This approach did not consider the current load on each coordinator, which led to an uneven distribution of connections and potential performance bottlenecks.

New behavior:

The impala-proxy now uses load-based routing to decide which coordinator should handle a new session request. The Impala proxy directs the new session to the coordinator with the minimum calculated load. You can customize how this load is calculated using the following parameters:

- `IMPALA_PROXY_COORDINATOR_LOAD_CPU_WEIGHT`: Determines the weight applied to the current percentage of CPU utilization when calculating the coordinator's load.
- `IMPALA_PROXY_COORDINATOR_LOAD_MEMORY_WEIGHT`: Determines the weight applied to the current percentage of memory utilization when calculating the coordinator's load.

By adjusting these weights, you can tune the Impala proxy to prioritize CPU or memory headroom when routing new sessions.

Summary:

New catalogd flag for HMS event sync defaults

Previous behavior:

Previously, disabling event processing globally required you to manually set the `impala.disableHmsSync` property for every individual database and table.

New behavior:

The new `disable_hms_sync_by_default` flag now defines the global default for event processing. If set to true, Impala skips event processing for all tables and databases unless the `impala.disableHmsSync` property is explicitly set to false at the table or database level. The priority for checking the sync status is the table property, followed by the database property, and finally the global default flag.

Apache Jira: [IMPALA-14085](#)

Behavioral Changes in Apache Iceberg

Behavioral changes denote a marked change in behavior from the previously released version to this version of Iceberg.

Cloudera Runtime 7.3.2

There are no behavioral changes in this release.

Behavioral Changes in Kafka

Functional adjustments and behavioral updates for Kafka are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Kafka, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see [Behavioral Changes](#).

Component-level custom Java home configuration removed

Previous behavior:

You could configure a component-specific Java home for Kafka Connect.

New behavior:

The component-level custom Java home configuration options are removed. Kafka Connect now uses the host-level `java_home` configuration. If you previously set a component-specific Java home for this service, verify the host-level `java_home` setting after upgrading.

The Kafka service is configured to use KRaft by default

Previous behavior:

Newly deployed Kafka services were configured to use ZooKeeper as their metadata store. The Kafka Metadata Store Service (`metadata.store`) Kafka property was set to ZooKeeper.

New behavior:

Newly deployed Kafka services are configured to use KRaft as their metadata store. The Kafka Metadata Store Service (`metadata.store`) Kafka property is set to KRaft.

High Watermark no longer advances when ISR is below MinISR

Previous behavior:

The High Watermark (HWM) advanced regardless of whether the in-sync replica (ISR) count was below `min.insync.replicas`. When producers used `acks=1` or `acks=0`, messages were written to the leader and became consumable once the HWM advanced, even if the ISR had dropped below the minimum threshold. The `min.insync.replicas` setting only affected `acks=all` produce requests, blocking writes when ISR was insufficient, but did not prevent HWM advancement or consumer reads for `acks=0/acks=1` messages.

New behavior:

The HWM no longer advances when the ISR count falls below `min.insync.replicas`. As a result, consumers are blocked from reading new messages in this condition, even if producers with `acks=1` or `acks=0` are still writing to the leader. This ensures data is only consumable when it meets the cluster's minimum durability requirements. If you use `min.insync.replicas=2` or higher, you may see reduced consumer throughput when the ISR count drops below the configured minimum.

For more information, see [KIP-966: Eligible Leader Replicas](#) and [KAFKA-15583](#).

Kafka protocol version is set automatically during upgrades

Previous behavior:

The `inter.broker.protocol.version` property for ZooKeeper-based clusters and the `metadata.version` property for KRaft-based clusters were not set automatically before an upgrade. Manually configuring these properties to the current protocol and metadata version was required before an upgrade.

New behavior:

During a cluster upgrade, Cloudera Manager now automatically sets the `inter.broker.protocol.version` property for ZooKeeper-based clusters and the `metadata.version` property for KRaft-based clusters. Manual configuration is no longer required.

**Note:**

Manually clearing these properties after the upgrade is still necessary. However, both properties are now available for direct configuration in Cloudera Manager. Using advanced configuration snippets is no longer required.

Behavioral Changes in Apache Knox

Behavioral changes denote a marked change in behavior from the previously released version to this version of Apache Knox.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Knox, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see [Behavioral Changes](#).

Summary:

The alias management configuration has been moved from service-level to role-level, with dedicated configurations for each role.

Previous behavior:

Previously, both Knox Gateway and IDBroker aliases were saved using the shared service-level `save_alias_command_input_password` configuration and a single Save Alias command.

New behavior:

Now, each role has its own dedicated role-level configuration: `gateway_save_alias_command_input` for the Knox Gateway role and `idbroker_save_alias_command_input` for the IDBroker role. Two role-specific commands are available: Save Alias - Knox Gateway and Save Alias - IDBroker. For more information, see [Saving aliases](#).

Behavioral Changes in Kudu

Functional adjustments and behavioral updates for Kudu are introduced in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

There are no behavioral changes in this release.

Behavioral Changes in Livy

Behavioral changes denote a marked change in behavior from the previously released version to this version of Livy.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Livy, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see [Livy Behavioral Changes](#).

There are no behavioral changes in this release.

Behavioral Changes in Oozie

Behavioral changes denote a marked change in behavior from the previously released version to this version of Oozie.

Cloudera Runtime 7.3.2

Oozie capture-output will not fail for long input data

Previous behavior:

Previously, the Oozie capture-output used to fail if the input data was longer than the maximum allowed length to be read.

New behavior:

Now, Oozie action works correctly even if the output data to be captured is longer than the maximum allowed length to be read. Instead, Oozie will read the first N (allowed by configuration) character from the output. To revert this behaviour, set `oozie.servlet.CallbackServlet.fail.on.data.overflow` and `oozie.action.fail.on.output.data.overflow` configurations to true in oozie-site safety valve.

Behavioral Changes in Ozone

Behavioral changes for Cloudera Runtime 7.3.2. introduce functional adjustments and behavioral updates for Ozone.

Cloudera Runtime 7.3.2

Migration of the `ozone debug recover` command to the `admin om lease recover` Ozone admin CLI command

The `ozone debug recover` command, used for lease recovery, now resides under the `admin om lease recover` Ozone admin CLI command.

Example,

```
# ozone admin om lease recover --path=ofs://<OZONE.SERVICE.ID>/
hsyncvol/hsyncbuck/hsync/File_0.txt
Lease recovery SUCCEEDED on ofs://<OZONE.SERVICE.ID>/hsyncvol/hs
yncbuck/hsync/File_0.txt
```

Changed functionality of the `--all` option in the `ozone sh snapshot lsDiff` command

Previous behavior:

Previously, the `--all` option was used to list snapshot diff jobs of any status.

New behavior:

Now, the `--all` option is used for listing limit. When the `--all` option is set, the listing limit is effectively set to `Integer.MAX_VALUE`. For the original `--all` behavior to list snapshot diff jobs, use the `--all-status` option instead.

Changed listing and deleting behavior at volume level

Previous behavior:

Previously, recursive listing and deleting a volume was supported using the `ozone fs` command that worked only if FSO bucket existed in volume.

New behavior:

Now, recursive listing and deleting a volume is supported only by using the `ozone shell` command and not the `ozone fs` command. The `ozone shell` command works for both OBS and FSO buckets inside a volume.

Example

- Recursive delete command

```
# ozone fs -rm -R -skipTrash ofs://ozone1713158544/vol/
```

Sample output

```
24/04/15 13:12:47 WARN ozone.BasicRootedOzoneFileSystem: Recursive volume delete using ofs is not supported
rm: Recursive volume delete using ofs is not supported. Instead use 'ozone sh volume delete -r -skipTrash -id <OM_SERVICE_ID> <Volume_URI>' command
```

- Recursive list command

```
# ozone fs -ls -R ofs://ozone1713158544/vol/
```

Sample output

```
24/04/15 13:12:53 WARN ozone.BasicRootedOzoneFileSystem: Recursive root/volume list using ofs is not supported
ls: Recursive list root/volume using ofs is not supported. Instead use 'ozone sh key list <Volume_URI>' command
```

Migration of the ozone debug chunkinfo API to the ozone debug replicas chunk-info API**Previous behavior:**

Previously, the ozone debug chunkinfo API returned the following JSON response:

```
{
  "KeyLocations" : [ [ {
    "Locations" : {
      "files" : [ "<block_path>" ],
      "pipelineID" : "<pipeline_id>"
    },
    "Datanode-HostName" : "<host_name>",
    "Datanode-IP" : "<host_ip>",
    "Container-ID" : <id>,
    "Block-ID" : <blockid>
  } ], { }, { }
  ] ] }
```

New behavior:

Now, the ozone debug chunkinfo API is relocated to the ozone debug replicas chunk-info API and returns the following JSON response:

```
{
  "volumeName" : "<vol>",
  "bucketName" : "<buck>",
  "name" : "<key>",
  "keyLocations" : [ [ {
    "datanode" : {
      "hostname" : "<host_name>",
      "ip" : "<host_ip>",
      "uuid" : "<uuid>"
    },
    "file" : "<block_path>",
    "blockData" : {
      "blockID" : {
        "containerID" : <id>,

```

```

    "localID" : <localid>,
    "blockCommitSequenceId" : <bcsid>
  },
  "size" : <size>,
  "chunks" : [ {
    "offset" : <offset>,
    "len" : <len>,
    "checksums" : [ "<checksum>" ],
    "checksumType" : "<ctype>",
    "bytesPerChecksum" : <bpc>
  } ]
},
{ }, { }
]]}

```

As part of this API update, the following keys are updated:

Table 8: JSON key mapping for the relocated ozone debug chunkinfo API

| Key name | Key in the previous JSON | Key in the current JSON |
|--------------|---|------------------------------------|
| Block path | data["KeyLocations"][0][0]["Locations"]["files"][0] | data["keyLocations"][0][0]["file"] |
| Container ID | "Container-ID" | blockData.blockID.containerID |
| Block ID | "Block-ID" | blockData.blockID.localID |
| Hostname | "Datanode-HostName" | datanode.hostname |

The S3 Gateway web UI is moved to a new port

Previous behavior:

The S3 Gateway web UI was accessible on ports 9878 for HTTP and 9879 for HTTPS, just like S3 buckets.

New behavior:

The S3 Gateway web UI resides on new ports to avoid conflicts with S3 buckets. The new ports are 19878 for HTTP and 19879 for HTTPS.

Behavioral Changes in Parquet

Functional adjustments and behavioral updates for Parquet are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

There are no behavioral changes in this release.

Behavioral Changes in Apache Ranger

Behavioral changes denote a marked change in behavior from the previously released version to this version of Apache Ranger.

Cloudera Runtime 7.3.2.0

The behavioral changes for Apache Ranger in Cloudera Runtime 7.3.2.0 include all cumulative updates from previous releases (such as 7.3.1.x). This version specifically incorporates changes introduced in Cloudera Runtime 7.3.1.100 through 7.3.1.706 alongside the following functional adjustments. For a complete list, see [Behavioral Changes](#).

Summary:

The new column authorization optimization property in the Ranger-HBase plugin changes Ranger audit behavior, when enabled. There is no behavioral change if the property is disabled.



Note: There may not be any performance increase if the user has explicit denials at the column level for any of the column families or the user is not fully authorized for the column families.

Previous behavior:

Earlier, there were more audit entries when the service configuration for optimization was not present and, hence, not enabled.

For example, the audit behavior for the `hbase ltt -tn multitest -families cf0,cf1 -write 3:2:1 -multiput -num_keys 1 -num_regions_per_server 1` command was as follows:

15 audit entries

New behavior:

Currently, there are fewer audit entries when the service configuration for optimization is enabled.

For example, the audit behavior for the `hbase ltt -tn multitest -families cf0,cf1 -write 3:2:1 -multiput -num_keys 1 -num_regions_per_server 1` command is as follows:

2 audit entries (multitest/cf1 and multitest/cf2)

Summary:

The following service configurations have been added to a new place in the Ranger Admin Web UI:

- Policy Download Users (policy.download.auth.users)
- Tag Download Users (tag.download.auth.users)
- Service Admin Users (service.admin.users)
- Service Admin Groups (service.admin.groups)
- Superusers (ranger.plugin.super.users)
- Superuser Groups (ranger.plugin.super.groups)
- Userstore Download Users (userstore.download.auth.users)

Previous behavior:

Previously, the above service configurations were available under the Add New Custom Configurations section.

New behavior:

Now, the above service configurations are directly available under the Config Properties section. Additionally, the configurations are added as dropdowns, where you can select the users and groups.

Also, after you upgrade to Cloudera Runtime 7.3.2.0 from any previous release, your existing configurations will be shifted to the Config Properties section with values.

Summary:

Policy resources in the Ranger Admin UI are being added using React JS instead of Backbone JS

Previous behavior:

Earlier with Backbone JS, when you copied and pasted resource values containing commas or spaces (for example, in Hive policy resources: database1, database2), the UI automatically split them into separate values — database1 and database2. The same behaviour applied to space-separated values. Because of this, you were not allowed to enter resource names containing commas and spaces, and this limitation affected all service policy resources.

New behavior:

After upgrading from Backbone JS to React JS, this restriction has been removed. Now, React JS treats pasted values with commas or spaces as a single entry. Hence, you can no longer paste multiple values at once; you must manually add each resource value.

Summary:

Hive authorization from Ranger for Alter Table Rename command does not require CREATE database permission on the database where the renamed table will be created.

Previous behavior:

Earlier, whenever Alter Table Rename command was used across databases in Hive, authorization from Ranger required CREATE database permission for the user on the target database in which the renamed table was created.

New behavior:

Now, whenever Alter Table Rename command is used across databases in Hive, authorization from Ranger does not check for CREATE database permission for the user on the target database in which the renamed table will be created.

Summary:

Added support for multiple columns policy creation in Ranger for Grant/Revoke request.

Previous behavior:

Previously, when a request with multiple columns, such as GRANT SELECT (col1, col2, col3, col4, col5, col6, col7, col8, col9, col10) ON TABLE demo.data5 TO ROLE testrole_09289898, is executed in Impala, it results in the creation of a separate grant policy for each column in Ranger.

New behavior:

Now, a request with multiple columns results in creation of a single policy for Grant request for all the columns in Ranger. Same is true for Revoke request.

Behavioral Changes in Ranger KMS

Behavioral changes denote a marked change in behavior from the previously released version to this version of Ranger KMS.

Cloudera Runtime 7.3.2.0

There are no behavioral changes introduced in this release.

Behavioral Changes in Schema Registry

Functional adjustments and behavioral updates for Schema Registry are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

Component-level custom Java home configuration removed

Previous behavior:

You could configure a component-specific Java home for Schema Registry.

New behavior:

The component-level custom Java home configuration options are removed. Schema Registry now uses the host-level java_home configuration. If you previously set a component-specific Java home for this service, verify the host-level java_home setting after upgrading.

Schema Registry now defaults to IPv4-only communication

The default value of the `schema.registry.additional.java.options` configuration parameter was updated to set the IP protocol to IPv4.

If you changed the default value of this parameter before upgrading, the new default value is not applied on upgrade. You can apply it manually after the upgrade.

Behavioral Changes in Sqoop

Functional adjustments and behavioral updates for Sqoop are introduced in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2

There are no behavioral changes in this release.

Behavioral Changes in Streams Messaging Manager

Functional adjustments and behavioral updates for Streams Messaging Manager are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

Component-level custom Java home configuration removed

Previous behavior:

You could configure a component-specific Java home for Streams Messaging Manager.

New behavior:

The component-level custom Java home configuration options are removed. Streams Messaging Manager now uses the host-level `java_home` configuration. If you previously set a component-specific Java home for this service, verify the host-level `java_home` setting after upgrading.

Default JMX settings changed to restrict connections to localhost

Previous behavior:

The default value of the `SMM_JMX_OPTS` Cloudera Manager configuration option was `-Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.authenticate=false -Dcom.sun.management.jmxremote.ssl=false`, which allowed unrestricted, unauthenticated JMX connections.

New behavior:

The default value of `SMM_JMX_OPTS` has been changed to restrict JMX connections to localhost only and enable SSL. If you previously customized `SMM_JMX_OPTS`, your custom value is preserved on upgrade. To revert to open JMX, update `SMM_JMX_OPTS` in Cloudera Manager.

Streams Messaging Manager now defaults to IPv4-only communication

A new argument was added to `SMM_JVM_PERF_OPTS` that sets the IP protocol to IPv4 by default.

If you changed the default value of this parameter before upgrading, the new default value is not applied on upgrade. You can apply it manually after the upgrade.

Streams Messaging Manager UI Migration to Java

The Streams Messaging Manager UI service is migrated from a NodeJS runtime to a Java-based server. This change addresses security vulnerabilities associated with NodeJS dependencies and aligns Streams Messaging Manager with the centralized dependency management of the platform.

As a result of this migration, the following changes apply:

- Runtime environment

The Streams Messaging Manager UI service now runs on the JVM. Configuration for the runtime environment is now managed via the `SMM_JAVA_OPTS` environment variable.
- TLS configuration

TLS configuration moved from OpenSSL-style parameters to standard Java JSSE configuration. New Cloudera Manager parameters manage TLS protocols and cipher suites:

 - `streams.messaging.manager.ui.ssl.supportedCipherSuites`
 - `streams.messaging.manager.ui.ssl.excludedCipherSuites`
 - `streams.messaging.manager.ui.ssl.supportedProtocols`
 - `streams.messaging.manager.ui.ssl.excludedProtocols`
- Configuration migration

During upgrade, Cloudera Manager attempts to automatically migrate existing TLS settings (including those found in the `NODE_OPTIONS` environment variable within safety valves) to the new Java-based configuration parameters. However, manual verification is strongly recommended.
- Safety valves

Any properties previously set in the Streams Messaging Manager UI Server Environment Advanced configuration Snippet (Safety Valve) using `NODE_OPTIONS` that are not related to TLS must be manually translated to their Java equivalents (if applicable) and set using `SMM_JAVA_OPTS`.

Behavioral Changes in Streams Replication Manager

Functional adjustments and behavioral updates for Streams Replication Manager are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

Component-level custom Java home configuration removed

Previous behavior:

You could configure a component-specific Java home for Streams Replication Manager.

New behavior:

The component-level custom Java home configuration options are removed. Streams Replication Manager now uses the host-level `java_home` configuration. If you previously set a component-specific Java home for this service, verify the host-level `java_home` setting after upgrading.

Default JMX settings changed to restrict connections to localhost

Previous behavior:

The default value of the `SRM_JMX_OPTS` Cloudera Manager configuration option was `-Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.authenticate=false -Dcom.sun.management.jmxremote.ssl=false`, which allowed unrestricted, unauthenticated JMX connections.

New behavior:

The default value of `SRM_JMX_OPTS` has been changed to restrict JMX connections to localhost only and enable SSL. If you previously customized `SRM_JMX_OPTS`, your custom value is preserved on upgrade. To revert to open JMX, update `SRM_JMX_OPTS` in Cloudera Manager.

Public Kafka Connect endpoints removed from SRM REST server

The public Kafka Connect endpoints are removed from the SRM REST server. Previously, these endpoints allowed users to interact with the internal Kafka Connect cluster inside the SRM instance,

including starting arbitrary connectors and modifying internal connectors. These endpoints were undocumented and not part of the official SRM API. They are no longer available.

Streams Replication Manager now defaults to IPv4-only communication

A new argument was added to `SRM_JVM_PERF_OPTS` that sets the IP protocol to IPv4 by default.

If you changed the default value of this parameter before upgrading, the new default value is not applied on upgrade. You can apply it manually after the upgrade.

Change in internal topic filtering logic

Summary:

The logic that identifies and filters internal topics in Streams Replication Manager has changed. This enables the replication of topics that appear to be internal but are not truly internal to Kafka and Streams Replication Manager, reducing the risk of unintentionally excluding user topics from replication.

Previous behavior:

Topics were filtered from replication by the `DefaultReplicationPolicy` and `IdentityReplicationPolicy` policies if the topic name ended with `[***SEPARATOR***]internal`. For example, `.internal`. In addition, the default deny list regex pattern was the following:

```
.*[\\-\\.].internal, .*\\.replica, __.*
```

New behavior:

Topics are now filtered from replication by the `DefaultReplicationPolicy` and `IdentityReplicationPolicy` policies if:

- Name starts with `mm2` and ends with `[***SEPARATOR***]internal`
- Name ends with `[***SEPARATOR***]checkpoints[***SEPARATOR***]internal`

In addition, the default deny list regex pattern is now the following:

```
mm2.*\\.internal, .*\\.replica, __.*
```

Internal Kafka topics that start with a dot (.) or two underscores (__) are continued to be filtered from replication by default.



Note: The `ReplicationPolicy` interface also changed. Previously, topics were filtered from replication if their names ended with `.internal` or `-internal`. From now on, names ending with `-internal` are no longer filtered. In addition, topics are filtered from replication if:

- Name starts with `mm2` and ends with `.internal`
- Name ends with `.checkpoints.internal`

Internal Kafka topics that start with a dot (.) or two underscores (__) are continued to be filtered from replication by default. If you have a custom policy that implements the interface, you might need to update it to correctly filter internal topics.

Behavioral Changes in Spark

Behavioral changes denote a marked change in behavior from the previously released version to this version of Spark.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Spark, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see [Spark Behavioral Changes](#).

Summary: Enhanced the security of Spark by implementing new default configuration settings**Previous behavior:**

1. `spark.ui.enabled=true`: possible initiation of an HTTP service that can be accessed from external hosts.
2. `spark.io.encryption.keySizeBits=128`: the default Spark `keySizeBits`

New behavior:

1. `spark.ui.enabled=false`: preventing initiation of an HTTP service that can be accessed from external hosts.
2. `spark.io.encryption.keySizeBits=256`: the default Spark `keySizeBits` has been increased from 128 to 256

Behavioral Changes in Spark Atlas Connector

Behavioral changes denote a marked change in behavior from the previously released version to this version of Spark Atlas Connector.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Spark Atlas Connector, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see [Spark Atlas Connector Behavioral Changes](#).

There are no behavioral changes in this release.

Behavioral Changes in YARN

Functional adjustments and behavioral updates for YARN and YARN Queue Manager are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2

There are no behavioral changes introduced in this release.

Behavioral Changes in ZooKeeper

Behavioral changes denote a marked change in behavior from the previously released version to this version of ZooKeeper.

Cloudera Runtime 7.3.2

ZooKeeper authentication enforcement

Previous behavior:

ZooKeeper client connections were accepted even when they did not authenticate. Anonymous clients could open sessions with the ZooKeeper ensemble, as long as they could reach the client port (for example, 2181). This meant that components and tools that did not use ZooKeeper authentication could still connect and operate, as long as ACLs and network controls (such as firewalls) allowed it.

New behavior:

When ZooKeeper authentication enforcement is enabled, the ZooKeeper server only accepts connections from clients that successfully authenticate (for example, via Kerberos/SASL). Components and clients that do not authenticate, or that fail authentication, can no longer establish a session with ZooKeeper and therefore cannot connect or operate through the unsecured client port.

This change hardens ZooKeeper security, but requires that all ZooKeeper#using components in the deployment be updated or configured to use authentication before enforcement is turned on.

Release Matrix

You must be familiar with the versions of all the service packs and cumulative hotfixes for Cloudera Runtime 7.3.2.

| Runtime Release Build Number | Cumulative Hotfix/Service Pack | Release Date | Download URL |
|------------------------------|--------------------------------|----------------|---|
| 7.3.2.0-957 | 7.3.2 | March 31, 2026 | https://archive.cloudera.com/p/cdh7/7.3.2/ |

Cloudera Runtime Component Versions

List of the official component versions for Cloudera Runtime. To know the component versions for compatibility with other applications, you must be familiar with the latest component versions in Cloudera Runtime. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Cloudera Runtime 7.3.2

List of the official component versions for Cloudera Runtime 7.3.2.

Apache Components

| Component | Version |
|--|-----------------------|
| Apache Arrow | 15.0.0.7.3.2.0-957 |
| Apache Atlas | 2.4.0.7.3.2.0-957 |
| Apache Calcite | 1.33.0.7.3.2.0-957 |
| Apache Avatica | 1.26.0.7.3.2.0-957 |
| Apache Avro | 1.11.3.7.3.2.0-957 |
| Apache Hadoop (Includes YARN and HDFS) | 3.4.2.7.3.2.0-957 |
| Apache HBase | 2.6.3.7.3.2.0-957 |
| Apache Hive | 3.1.3000.7.3.2.0-957 |
| Apache Iceberg | 1.5.2.7.3.2.0-957 |
| Apache Impala | 4.5.0.7.3.2.0-957 |
| Apache Kafka | 3.9.1.7.3.2.0-957 |
| Apache Knox | 2.1.0.7.3.2.0-957 |
| Apache Kudu | 1.18.0.7.3.2.0-957 |
| Apache Livy | 0.7.23000.7.3.2.0-957 |
| Apache Ozone | 2.0.0.7.3.2.0-957 |
| Apache Oozie | 5.1.0.7.3.2.0-957 |
| Apache ORC | 2.0.3.7.3.2.0-957 |
| Apache Parquet | 1.15.2.7.3.2.0-957 |
| Apache Phoenix | 5.2.1.7.3.2.0-957 |

| Component | Version |
|------------------|--------------------|
| Apache Ranger | 2.6.0.7.3.2.0-957 |
| Apache Solr | 8.11.2.7.3.2.0-957 |
| Apache Spark | 3.5.4.7.3.2.0-957 |
| Apache Sqoop | 1.4.7.7.3.2.0-957 |
| Apache Tez | 0.9.1.7.3.2.0-957 |
| Apache ZooKeeper | 3.8.5.7.3.2.0-957 |

Other Components

| Component | Version |
|-----------------------------|---------------------|
| Cruise Control | 2.5.143.7.3.2.0-957 |
| Hue | 4.11.0.7.3.2.0-957 |
| Search | 1.0.0.7.3.2.0-957 |
| Schema Registry | 0.10.0.7.3.2.0-957 |
| Streams Messaging Manager | 2.3.0.7.3.2.0-957 |
| Streams Replication Manager | 1.1.0.7.3.2.0-957 |

Connectors and Encryption Components

| Component | Version |
|--------------------------|-------------------|
| HBase connectors | 1.0.0.7.3.2.0-957 |
| Hive Warehouse Connector | 1.0.0.7.3.2.0-957 |
| Spark Atlas Connector | 3.5.4.7.3.2.0-957 |
| Spark Schema Registry | 3.5.4.7.3.2.0-957 |

Using the Cloudera Runtime Maven repository

Information about using Maven to build applications with Cloudera Runtime components.

If you want to build applications or tools for use with Cloudera Runtime components and you are using Maven or Ivy for dependency management, you can pull the Cloudera Runtime artifacts from the Cloudera Maven repository. The repository is available at <https://repository.cloudera.com/artifactory/cloudera-repos/>.



Important: When you build an application JAR, do not include CDH JARs, because they are already provided. If you do, upgrading CDH can break your application. To avoid this situation, set the Maven dependency scope to provided. If you have already built applications which include the CDH JARs, update the dependency to set scope to provided and recompile.

The following is a sample POM (pom.xml) file:

```
<project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/maven-v4_0_0.xsd">
  <repositories>
    <repository>
      <id>cloudera</id>
      <url>https://repository.cloudera.com/artifactory/cloudera-repos/</url>
    </repository>
  </repositories>
```

```
</project>
```

Cloudera Runtime 7.3.2.0-957

The following table lists the project name, groupId, artifactId, and version required to access each RUNTIME artifact.

| Project | groupId | artifactId | version |
|---------|------------------|------------------------------|-------------------|
| Atlas | org.apache.atlas | atlas-authorization | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-aws-s3-bridge | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-azure-adls-bridge | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-classification-updater | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-client-common | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-client-v1 | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-client-v2 | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-client-v2-shaded | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-common | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-distro | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-docs | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-graphdb-api | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-graphdb-common | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-graphdb-janus | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-hdfs-bridge | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-index-repair-tool | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-intg | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-notification | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-notification-analyzer | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-plugin-classloader | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-repository | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-server-api | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | atlas-testtools | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | hbase-bridge | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | hbase-bridge-shim | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | hbase-testing-util | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | hdfs-model | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | hive-bridge | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | hive-bridge-shim | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | impala-bridge | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | impala-bridge-shim | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | impala-hook-api | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | kafka-bridge | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | kafka-bridge-shim | 2.4.0.7.3.2.0-957 |

| Project | groupId | artifactId | version |
|---------|--------------------|------------------------------|--------------------|
| Atlas | org.apache.atlas | navigator-to-atlas | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | sample-app | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | sqoop-bridge | 2.4.0.7.3.2.0-957 |
| Atlas | org.apache.atlas | sqoop-bridge-shim | 2.4.0.7.3.2.0-957 |
| Avro | org.apache.avro | avro | 1.11.3.7.3.2.0-957 |
| Avro | org.apache.avro | avro-android | 1.11.3.7.3.2.0-957 |
| Avro | org.apache.avro | avro-codegen-test | 1.11.3.7.3.2.0-957 |
| Avro | org.apache.avro | avro-compiler | 1.11.3.7.3.2.0-957 |
| Avro | org.apache.avro | avro-grpc | 1.11.3.7.3.2.0-957 |
| Avro | org.apache.avro | avro-ipc | 1.11.3.7.3.2.0-957 |
| Avro | org.apache.avro | avro-ipc-jetty | 1.11.3.7.3.2.0-957 |
| Avro | org.apache.avro | avro-ipc-netty | 1.11.3.7.3.2.0-957 |
| Avro | org.apache.avro | avro-mapred | 1.11.3.7.3.2.0-957 |
| Avro | org.apache.avro | avro-maven-plugin | 1.11.3.7.3.2.0-957 |
| Avro | org.apache.avro | avro-perf | 1.11.3.7.3.2.0-957 |
| Avro | org.apache.avro | avro-protobuf | 1.11.3.7.3.2.0-957 |
| Avro | org.apache.avro | avro-service-archetype | 1.11.3.7.3.2.0-957 |
| Avro | org.apache.avro | avro-test-custom-conversions | 1.11.3.7.3.2.0-957 |
| Avro | org.apache.avro | avro-thrift | 1.11.3.7.3.2.0-957 |
| Avro | org.apache.avro | avro-tools | 1.11.3.7.3.2.0-957 |
| Avro | org.apache.avro | trevni-avro | 1.11.3.7.3.2.0-957 |
| Avro | org.apache.avro | trevni-core | 1.11.3.7.3.2.0-957 |
| Calcite | org.apache.calcite | calcite-babel | 1.33.0.7.3.2.0-957 |
| Calcite | org.apache.calcite | calcite-core | 1.33.0.7.3.2.0-957 |
| Calcite | org.apache.calcite | calcite-druid | 1.33.0.7.3.2.0-957 |
| Calcite | org.apache.calcite | calcite-kafka | 1.33.0.7.3.2.0-957 |
| Calcite | org.apache.calcite | calcite-linq4j | 1.33.0.7.3.2.0-957 |
| Calcite | org.apache.calcite | calcite-plus | 1.33.0.7.3.2.0-957 |
| Calcite | org.apache.calcite | calcite-server | 1.33.0.7.3.2.0-957 |
| Calcite | org.apache.calcite | calcite-testkit | 1.33.0.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-aliyun | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-annotations | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-archive-logs | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-archives | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-assemblies | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-auth | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-aws | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-azure | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-azure-datalake | 3.4.2.7.3.2.0-957 |

| Project | groupId | artifactId | version |
|---------|-------------------|------------------------------------|-------------------|
| Hadoop | org.apache.hadoop | hadoop-benchmark | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-build-tools | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-client | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-client-api | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-client-integration-tests | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-client-minicluster | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-client-runtime | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-cloud-storage | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-common | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-datajoin | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-distcp | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-dynamometer-blockgen | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-dynamometer-dist | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-dynamometer-infra | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-dynamometer-workload | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-extras | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-federation-balance | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-fs2img | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-gridmix | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-hdfs | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-hdfs-client | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-hdfs-httpfs | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-hdfs-native-client | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-hdfs-nfs | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-hdfs-rbf | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-kafka | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-kms | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-mapreduce-client-app | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-mapreduce-client-common | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-mapreduce-client-core | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-mapreduce-client-hs | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-mapreduce-client-hs-plugins | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-mapreduce-client-jobclient | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-mapreduce-client-nativetask | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-mapreduce-client-shuffle | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-mapreduce-client-uploader | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-mapreduce-examples | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-maven-plugins | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-minicluster | 3.4.2.7.3.2.0-957 |

| Project | groupId | artifactId | version |
|---------|-------------------|---|-------------------|
| Hadoop | org.apache.hadoop | hadoop-minikdc | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-nfs | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-openstack | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-registry | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-resourceestimator | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-rumen | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-sls | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-streaming | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-tools-dist | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-api | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-applications-distributedshell | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-applications-unmanaged-am-launcher | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-client | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-common | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-csi | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-registry | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-server-applicationhistoryservice | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-server-common | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-server-globalpolicygenerator | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-server-nodemanager | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-server-resourcemanager | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-server-router | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-server-sharedcachemanager | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-server-tests | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-server-timeline-pluginstorage | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-server-timelineservice | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-server-timelineservice-documentstore | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-client | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-common | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-server-2 | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-tests | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-server-web-proxy | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-services-api | 3.4.2.7.3.2.0-957 |
| Hadoop | org.apache.hadoop | hadoop-yarn-services-core | 3.4.2.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-annotations | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-asyncfs | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-backup | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-checkstyle | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-client | 2.6.3.7.3.2.0-957 |

| Project | groupId | artifactId | version |
|---------|------------------|----------------------------------|----------------------|
| HBase | org.apache.hbase | hbase-client-project | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-common | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-compression-aircompressor | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-compression-brotli | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-compression-lz4 | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-compression-snappy | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-compression-zstd | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-endpoint | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-examples | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-external-blockcache | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-hadoop-compat | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-hadoop2-compat | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-hbtop | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-http | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-it | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-logging | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-mapreduce | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-metrics | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-metrics-api | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-openssl | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-procedure | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-protocol | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-protocol-shaded | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-replication | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-resource-bundle | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-rest | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-rsgroup | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-server | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-shaded-client | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-shaded-client-byo-hadoop | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-shaded-client-project | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-shaded-mapreduce | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-shaded-testing-util | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-shaded-testing-util-tester | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-shell | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-testing-util | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-thrift | 2.6.3.7.3.2.0-957 |
| HBase | org.apache.hbase | hbase-zookeeper | 2.6.3.7.3.2.0-957 |
| Hive | org.apache.hive | catalogd-unit | 3.1.3000.7.3.2.0-957 |

| Project | groupId | artifactId | version |
|---------|-----------------|-----------------------|----------------------|
| Hive | org.apache.hive | hive-beeline | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-blobstore | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-classification | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-cli | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-common | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-contrib | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-exec | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-hbase-handler | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-hcatalog-it-unit | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-hms-catalog | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-hplsql | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-iceberg-catalog | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-iceberg-handler | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-iceberg-shading | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-impala | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-it-custom-serde | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-it-iceberg | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-it-impala | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-it-minikdc | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-it-qfile | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-it-qfile-kudu | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-it-test-serde | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-it-unit | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-it-unit-hadoop2 | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-it-util | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-jdbc | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-jdbc-handler | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-jmh | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-kudu-handler | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-llap-client | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-llap-common | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-llap-ext-client | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-llap-server | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-llap-tez | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-metastore | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-parser | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-serde | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-service | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-service-rpc | 3.1.3000.7.3.2.0-957 |

| Project | groupId | artifactId | version |
|--------------------------|----------------------|--|--------------------------------------|
| Hive | org.apache.hive | hive-shims | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-standalone-metastore | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-storage-api | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-streaming | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-testutils | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-udf | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | hive-vector-code-gen | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | kafka-handler | 3.1.3000.7.3.2.0-957 |
| Hive | org.apache.hive | patched-iceberg-api | patched-1.5.2.7.3.2.0-957-3.1.3000.7 |
| Hive | org.apache.hive | patched-iceberg-core | patched-1.5.2.7.3.2.0-957-3.1.3000.7 |
| Hive Warehouse Connector | com.hortonworks.hive | hive-warehouse-connector-spark3_2.12 | 1.0.0.7.3.2.0-957 |
| Kafka | org.apache.kafka | connect | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | connect-api | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | connect-basic-auth-extension | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | connect-cloudera-authorization-extension | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | connect-cloudera-common | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | connect-cloudera-secret-storage | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | connect-cloudera-security-policies | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | connect-file | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | connect-json | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | connect-mirror | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | connect-mirror-client | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | connect-runtime | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | connect-test-plugins | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | connect-transforms | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | generator | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-clients | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-cloudera-metrics-reporter_2.12 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-cloudera-metrics-reporter_2.13 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-cloudera-plugins | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-examples | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-group-coordinator | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-group-coordinator-api | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-log4j-appender | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-metadata | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-raft | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-server | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-server-common | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-shell | 3.9.1.7.3.2.0-957 |

| Project | groupId | artifactId | version |
|---------|------------------|---|-------------------|
| Kafka | org.apache.kafka | kafka-storage | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-storage-api | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-examples | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-scala_2.12 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-scala_2.13 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-test-utils | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-0100 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-0101 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-0102 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-0110 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-10 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-11 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-20 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-21 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-22 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-23 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-24 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-25 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-26 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-27 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-28 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-30 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-31 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-32 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-33 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-34 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-35 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-36 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-37 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-streams-upgrade-system-tests-38 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-tools | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-tools-api | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka-transaction-coordinator | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka_2.12 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | kafka_2.13 | 3.9.1.7.3.2.0-957 |
| Kafka | org.apache.kafka | trogdor | 3.9.1.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-adapter | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-admin-ui | 2.1.0.7.3.2.0-957 |

| Project | groupId | artifactId | version |
|---------|-----------------|---|-------------------|
| Knox | org.apache.knox | gateway-applications | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-cloud-bindings | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-demo-ldap | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-demo-ldap-launcher | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-discovery-ambari | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-discovery-cm | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-docker | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-i18n | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-i18n-logging-log4j | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-i18n-logging-slf4j | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-openapi-ui | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-performance-test | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-ha | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-identity-assertion-common | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-identity-assertion-concat | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-identity-assertion-hadoop-groups | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-identity-assertion-no-doas | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-identity-assertion-pseudo | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-identity-assertion-regex | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-identity-assertion-switchcase | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-jersey | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-rewrite | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-rewrite-common | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-rewrite-func-hostmap-static | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-rewrite-func-inbound-query-param | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-rewrite-func-service-registry | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-rewrite-step-encrypt-uri | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-rewrite-step-secure-query | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-security-authc-anon | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-security-authc-remote | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-security-authz-acls | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-security-authz-composite | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-security-authz-path-acls | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-security-clientcert | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-security-hadoopauth | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-security-jwt | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-security-pac4j | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-security-preauth | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-provider-security-shiro | 2.1.0.7.3.2.0-957 |

| Project | groupId | artifactId | version |
|---------|-----------------|-------------------------------------|-------------------|
| Knox | org.apache.knox | gateway-provider-security-webappsec | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-release | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-server | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-server-launcher | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-server-xforwarded-filter | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-admin | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-as | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-auth | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-definitions | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-hashicorp-vault | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-hbase | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-health | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-hive | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-idbroker | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-idbroker-plugins | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-impala | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-jkg | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-knoxssso | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-knoxssout | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-knoxtoken | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-livy | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-metadata | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-nifi | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-nifi-registry | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-remoteconfig | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-rm | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-session | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-storm | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-test | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-tgs | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-vault | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-service-webhdfs | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-shell | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-shell-launcher | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-shell-release | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-shell-samples | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-spi | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-spi-common | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-test | 2.1.0.7.3.2.0-957 |

| Project | groupId | artifactId | version |
|---------|-----------------|-----------------------------|-----------------------|
| Knox | org.apache.knox | gateway-test-idbroker | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-test-release-utils | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-test-utils | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-topology-hadoop-xml | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-topology-simple | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-util-common | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-util-configinjector | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-util-launcher | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | gateway-util-urltemplate | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | hadoop-examples | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | knox-cli-launcher | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | knox-homepage-ui | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | knox-token-generation-ui | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | knox-token-management-ui | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | knox-webshell-ui | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | webhdfs-kerb-test | 2.1.0.7.3.2.0-957 |
| Knox | org.apache.knox | webhdfs-test | 2.1.0.7.3.2.0-957 |
| Kudu | org.apache.kudu | kudu-backup-common | 1.18.0.7.3.2.0-957 |
| Kudu | org.apache.kudu | kudu-backup-tools | 1.18.0.7.3.2.0-957 |
| Kudu | org.apache.kudu | kudu-backup3_2.12 | 1.18.0.7.3.2.0-957 |
| Kudu | org.apache.kudu | kudu-client | 1.18.0.7.3.2.0-957 |
| Kudu | org.apache.kudu | kudu-flatbuffers | 1.18.0.7.3.2.0-957 |
| Kudu | org.apache.kudu | kudu-hive | 1.18.0.7.3.2.0-957 |
| Kudu | org.apache.kudu | kudu-proto | 1.18.0.7.3.2.0-957 |
| Kudu | org.apache.kudu | kudu-replication | 1.18.0.7.3.2.0-957 |
| Kudu | org.apache.kudu | kudu-spark3-tools_2.12 | 1.18.0.7.3.2.0-957 |
| Kudu | org.apache.kudu | kudu-spark3_2.12 | 1.18.0.7.3.2.0-957 |
| Kudu | org.apache.kudu | kudu-subprocess | 1.18.0.7.3.2.0-957 |
| Kudu | org.apache.kudu | kudu-test-utils | 1.18.0.7.3.2.0-957 |
| Livy | org.apache.livy | livy-api | 0.7.23000.7.3.2.0-957 |
| Livy | org.apache.livy | livy-client-common | 0.7.23000.7.3.2.0-957 |
| Livy | org.apache.livy | livy-client-http | 0.7.23000.7.3.2.0-957 |
| Livy | org.apache.livy | livy-core_2.12 | 0.7.23000.7.3.2.0-957 |
| Livy | org.apache.livy | livy-examples | 0.7.23000.7.3.2.0-957 |
| Livy | org.apache.livy | livy-integration-test | 0.7.23000.7.3.2.0-957 |
| Livy | org.apache.livy | livy-repl_2.12 | 0.7.23000.7.3.2.0-957 |
| Livy | org.apache.livy | livy-rsc | 0.7.23000.7.3.2.0-957 |
| Livy | org.apache.livy | livy-scala-api_2.12 | 0.7.23000.7.3.2.0-957 |
| Livy | org.apache.livy | livy-server | 0.7.23000.7.3.2.0-957 |

| Project | groupId | artifactId | version |
|---------|-------------------|-----------------------------|-----------------------|
| Livy | org.apache.livy | livy-test-lib | 0.7.23000.7.3.2.0-957 |
| Livy | org.apache.livy | livy-thriftserver | 0.7.23000.7.3.2.0-957 |
| Livy | org.apache.livy | livy-thriftserver-session | 0.7.23000.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-analyzers-common | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-analyzers-icu | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-analyzers-kuromoji | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-analyzers-morfologik | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-analyzers-nori | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-analyzers-openslp | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-analyzers-phonetic | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-analyzers-smarten | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-analyzers-stempel | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-backward-codecs | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-benchmark | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-classification | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-codecs | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-core | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-demo | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-expressions | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-facet | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-grouping | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-highlighter | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-join | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-memory | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-misc | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-monitor | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-queries | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-queryparser | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-replicator | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-sandbox | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-spatial-extras | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-spatial3d | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-suggest | 8.11.2.7.3.2.0-957 |
| Lucene | org.apache.lucene | lucene-test-framework | 8.11.2.7.3.2.0-957 |
| Oozie | org.apache.oozie | oozie-client | 5.1.0.7.3.2.0-957 |
| Oozie | org.apache.oozie | oozie-core | 5.1.0.7.3.2.0-957 |
| Oozie | org.apache.oozie | oozie-distro | 5.1.0.7.3.2.0-957 |
| Oozie | org.apache.oozie | oozie-examples | 5.1.0.7.3.2.0-957 |
| Oozie | org.apache.oozie | oozie-fluent-job-api | 5.1.0.7.3.2.0-957 |

| Project | groupId | artifactId | version |
|---------|------------------|--------------------------------|-------------------|
| Oozie | org.apache.oozie | oozie-fluent-job-client | 5.1.0.7.3.2.0-957 |
| Oozie | org.apache.oozie | oozie-server | 5.1.0.7.3.2.0-957 |
| Oozie | org.apache.oozie | oozie-sharelib-distcp | 5.1.0.7.3.2.0-957 |
| Oozie | org.apache.oozie | oozie-sharelib-git | 5.1.0.7.3.2.0-957 |
| Oozie | org.apache.oozie | oozie-sharelib-hcatalog | 5.1.0.7.3.2.0-957 |
| Oozie | org.apache.oozie | oozie-sharelib-hive | 5.1.0.7.3.2.0-957 |
| Oozie | org.apache.oozie | oozie-sharelib-hive2 | 5.1.0.7.3.2.0-957 |
| Oozie | org.apache.oozie | oozie-sharelib-oozie | 5.1.0.7.3.2.0-957 |
| Oozie | org.apache.oozie | oozie-sharelib-spark3 | 5.1.0.7.3.2.0-957 |
| Oozie | org.apache.oozie | oozie-sharelib-sqoop | 5.1.0.7.3.2.0-957 |
| Oozie | org.apache.oozie | oozie-sharelib-streaming | 5.1.0.7.3.2.0-957 |
| Oozie | org.apache.oozie | oozie-tools | 5.1.0.7.3.2.0-957 |
| Oozie | org.apache.oozie | oozie-zookeeper-security-tests | 5.1.0.7.3.2.0-957 |
| ORC | org.apache.orc | orc-core | 2.0.3.7.3.2.0-957 |
| ORC | org.apache.orc | orc-mapreduce | 2.0.3.7.3.2.0-957 |
| ORC | org.apache.orc | orc-shims | 2.0.3.7.3.2.0-957 |
| ORC | org.apache.orc | orc-tools | 2.0.3.7.3.2.0-957 |
| Ozone | org.apache.ozone | hdds-annotation-processing | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | hdds-client | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | hdds-common | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | hdds-config | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | hdds-container-service | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | hdds-crypto-api | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | hdds-crypto-default | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | hdds-docs | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | hdds-erasurencode | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | hdds-interface-admin | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | hdds-interface-client | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | hdds-interface-server | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | hdds-managed-rocksdbs | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | hdds-rocks-native | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | hdds-server-framework | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | hdds-server-scm | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | hdds-test-utils | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | mini-chaos-tests | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-cli-admin | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-cli-shell | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-client | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-common | 2.0.0.7.3.2.0-957 |

| Project | groupId | artifactId | version |
|---------|--------------------|------------------------------|--------------------|
| Ozone | org.apache.ozone | ozone-csi | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-datanode | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-dev-support | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-dist | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-filessystem | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-filessystem-common | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-filessystem-hadoop2 | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-filessystem-hadoop3 | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-filessystem-shaded | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-freon | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-httpfgateway | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-insight | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-integration-test | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-integration-test-recon | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-integration-test-s3 | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-interface-client | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-interface-storage | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-manager | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-mini-cluster | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-network-tests | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-recon | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-reconcodegen | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-s3-secret-store | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-s3gateway | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | ozone-tools | 2.0.0.7.3.2.0-957 |
| Ozone | org.apache.ozone | rocksdb-checkpoint-differ | 2.0.0.7.3.2.0-957 |
| Parquet | org.apache.parquet | parquet-avro | 1.15.2.7.3.2.0-957 |
| Parquet | org.apache.parquet | parquet-column | 1.15.2.7.3.2.0-957 |
| Parquet | org.apache.parquet | parquet-common | 1.15.2.7.3.2.0-957 |
| Parquet | org.apache.parquet | parquet-encoding | 1.15.2.7.3.2.0-957 |
| Parquet | org.apache.parquet | parquet-format-structures | 1.15.2.7.3.2.0-957 |
| Parquet | org.apache.parquet | parquet-generator | 1.15.2.7.3.2.0-957 |
| Parquet | org.apache.parquet | parquet-hadoop | 1.15.2.7.3.2.0-957 |
| Parquet | org.apache.parquet | parquet-hadoop-bundle | 1.15.2.7.3.2.0-957 |
| Parquet | org.apache.parquet | parquet-jackson | 1.15.2.7.3.2.0-957 |
| Parquet | org.apache.parquet | parquet-pig | 1.15.2.7.3.2.0-957 |
| Parquet | org.apache.parquet | parquet-pig-bundle | 1.15.2.7.3.2.0-957 |
| Parquet | org.apache.parquet | parquet-protobuf | 1.15.2.7.3.2.0-957 |
| Parquet | org.apache.parquet | parquet-thrift | 1.15.2.7.3.2.0-957 |

| Project | groupId | artifactId | version |
|---------|--------------------|--|--------------------|
| Parquet | org.apache.parquet | parquet-tools | 1.15.2.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix-client-embedded-hbase-2.6 | 5.2.1.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix-client-lite-hbase-2.6 | 5.2.1.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix-core | 5.2.1.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix-core-client | 5.2.1.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix-core-server | 5.2.1.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix-hbase-compat-2.4.1 | 5.2.1.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix-hbase-compat-2.5.0 | 5.2.1.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix-hbase-compat-2.5.4 | 5.2.1.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix-hbase-compat-2.6.0 | 5.2.1.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix-mapreduce-byo-shaded-hbase-hbase-2.6 | 5.2.1.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix-pherf | 5.2.1.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix-queryserver | 6.0.0.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-client | 6.0.0.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-it | 6.0.0.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-load-balancer | 6.0.0.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-orchestrator | 6.0.0.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix-server-hbase-2.6 | 5.2.1.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix-tracing-webapp | 5.2.1.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix5-hive | 6.0.0.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix5-hive-shaded | 6.0.0.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix5-spark3 | 6.0.0.7.3.2.0-957 |
| Phoenix | org.apache.phoenix | phoenix5-spark3-shaded | 6.0.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | conditions-enrichers | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | credentialbuilder | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | embeddedwebservice | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | hdfs-nn-perf | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | jisql | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ldapconfigcheck | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-adls-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-atlas-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-atlas-plugin-shim | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-authn | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-caii-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-clo-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-common-ha | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-distro | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-examples-distro | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-gs-plugin | 2.6.0.7.3.2.0-957 |

| Project | groupId | artifactId | version |
|---------|-------------------|-------------------------------|-------------------|
| Ranger | org.apache.ranger | ranger-hbase-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-hbase-plugin-shim | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-hdfs-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-hdfs-plugin-shim | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-hive-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-hive-plugin-shim | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-intg | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-kafka-connect-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-kafka-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-kafka-plugin-shim | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-kms | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-kms-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-kms-plugin-shim | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-knox-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-knox-plugin-shim | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-kudu-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-kylin-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-kylin-plugin-shim | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-metrics | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-nestedstructure-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-nifi-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-nifi-registry-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-opensearch-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-ozone-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-ozone-plugin-shim | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-plugin-classloader | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-plugin-perf | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-plugins-audit | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-plugins-common | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-plugins-cred | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-plugins-installer | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-policymigration | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-raz-adls | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-raz-chained-plugins | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-raz-hook-abfs | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-raz-hook-s3 | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-raz-intg | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-raz-processor | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-raz-s3 | 2.6.0.7.3.2.0-957 |

| Project | groupId | artifactId | version |
|---------|-------------------|--------------------------------|--------------------|
| Ranger | org.apache.ranger | ranger-raz-s3-lib | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-rms-common | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-rms-hive | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-rms-plugins-common | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-rms-tools | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-rms-webapp | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-s3-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-sampleapp-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-schema-registry-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-solr-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-solr-plugin-shim | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-sqoop-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-sqoop-plugin-shim | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-storm-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-storm-plugin-shim | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-tagsync | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-tools | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-trino-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-util | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-yarn-plugin | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ranger-yarn-plugin-shim | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | sample-client | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | sampleapp | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | ugsync-util | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | unixauthclient | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | unixauthservice | 2.6.0.7.3.2.0-957 |
| Ranger | org.apache.ranger | unixusersync | 2.6.0.7.3.2.0-957 |
| Solr | org.apache.solr | solr-analysis-extras | 8.11.2.7.3.2.0-957 |
| Solr | org.apache.solr | solr-analytics | 8.11.2.7.3.2.0-957 |
| Solr | org.apache.solr | solr-cell | 8.11.2.7.3.2.0-957 |
| Solr | org.apache.solr | solr-core | 8.11.2.7.3.2.0-957 |
| Solr | org.apache.solr | solr-dataimporthandler | 8.11.2.7.3.2.0-957 |
| Solr | org.apache.solr | solr-dataimporthandler-extras | 8.11.2.7.3.2.0-957 |
| Solr | org.apache.solr | solr-gcs-repository | 8.11.2.7.3.2.0-957 |
| Solr | org.apache.solr | solr-jaegertracer-configurator | 8.11.2.7.3.2.0-957 |
| Solr | org.apache.solr | solr-langid | 8.11.2.7.3.2.0-957 |
| Solr | org.apache.solr | solr-ltr | 8.11.2.7.3.2.0-957 |
| Solr | org.apache.solr | solr-prometheus-exporter | 8.11.2.7.3.2.0-957 |
| Solr | org.apache.solr | solr-s3-repository | 8.11.2.7.3.2.0-957 |

| Project | groupId | artifactId | version |
|---------|------------------|--|--------------------|
| Solr | org.apache.solr | solr-security-util | 8.11.2.7.3.2.0-957 |
| Solr | org.apache.solr | solr-solrj | 8.11.2.7.3.2.0-957 |
| Solr | org.apache.solr | solr-test-framework | 8.11.2.7.3.2.0-957 |
| Solr | org.apache.solr | solr-velocity | 8.11.2.7.3.2.0-957 |
| Spark | org.apache.spark | spark-avro_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-catalyst_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-common-utils_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-connect-client-jvm_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-connect-common_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-connect_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-core_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-graphx_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-hadoop-cloud_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-hive_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-kubernetes_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-kvstore_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-launcher_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-mllib-local_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-mllib_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-network-common_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-network-shuffle_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-network-yarn_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-protobuf_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-repl_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-shaded-raz | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-sketch_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-sql-api_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-sql-kafka-0-10_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-sql_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-streaming-kafka-0-10-assembly_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-streaming-kafka-0-10_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-streaming_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-tags_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-token-provider-kafka-0-10_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-unsafe_2.12 | 3.5.4.7.3.2.0-957 |
| Spark | org.apache.spark | spark-yarn_2.12 | 3.5.4.7.3.2.0-957 |
| Sqoop | org.apache.sqoop | sqoop | 1.4.7.7.3.2.0-957 |
| Sqoop | org.apache.sqoop | sqoop-test | 1.4.7.7.3.2.0-957 |
| Tez | org.apache.tez | hadoop-shim | 0.9.1.7.3.2.0-957 |

| Project | groupId | artifactId | version |
|-----------|----------------------|-------------------------------------|-------------------|
| Tez | org.apache.tez | hadoop-shim-2.8 | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-api | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-aux-services | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-build-tools | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-common | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-dag | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-examples | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-ext-service-tests | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-history-parser | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-javadoc-tools | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-job-analyzer | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-mapreduce | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-protobuf-history-plugin | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-runtime-internals | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-runtime-library | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-tests | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-yarn-timeline-cache-plugin | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-yarn-timeline-history | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-yarn-timeline-history-with-acls | 0.9.1.7.3.2.0-957 |
| Tez | org.apache.tez | tez-yarn-timeline-history-with-fs | 0.9.1.7.3.2.0-957 |
| ZooKeeper | org.apache.zookeeper | zookeeper | 3.8.5.7.3.2.0-957 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-fatjar | 3.8.5.7.3.2.0-957 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-loggraph | 3.8.5.7.3.2.0-957 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-rest | 3.8.5.7.3.2.0-957 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-zooinspector | 3.8.5.7.3.2.0-957 |
| ZooKeeper | org.apache.zookeeper | zookeeper-it | 3.8.5.7.3.2.0-957 |
| ZooKeeper | org.apache.zookeeper | zookeeper-jute | 3.8.5.7.3.2.0-957 |
| ZooKeeper | org.apache.zookeeper | zookeeper-prometheus-metrics | 3.8.5.7.3.2.0-957 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-election | 3.8.5.7.3.2.0-957 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-lock | 3.8.5.7.3.2.0-957 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-queue | 3.8.5.7.3.2.0-957 |

Fixed Common Vulnerabilities and Exposures

Common vulnerabilities and Exposures (CVEs) fixed in Cloudera Runtime 7.3.2 release, and its Service Packs and Cumulative hotfixes.

Cloudera Runtime 7.3.2

Common vulnerabilities and Exposures (CVEs) fixed in Cloudera Runtime 7.3.2.

| CVE ID | Description |
|----------------------------------|--|
| CVE-2016-1000027 | Pivotal Spring Framework through 5.3.16 suffers from a potential remote code execution (RCE) issue if used for Java deserialization of untrusted data. Depending on how the library is implemented within a product, this issue may or not occur, and authentication may be required. NOTE: the vendor's position is that untrusted data is not an intended use case. The product's behavior will not be changed because some users rely on deserialization of trusted data. |
| CVE-2021-4048 | An out-of-bounds read flaw was found in the CLARRV, DLARRV, SLARRV, and ZLARRV functions in lapack through version 3.10.0, as also used in OpenBLAS before version 0.3.18. Specially crafted inputs passed to these functions could cause an application using lapack to crash or possibly disclose portions of its memory. |
| CVE-2024-38821 | Spring WebFlux applications that have Spring Security authorization rules on static resources can be bypassed under certain circumstances. For this to impact an application, all of the following must be true: * It must be a WebFlux application * It must be using Spring's static resources support * It must have a non-permitAll authorization rule applied to the static resources support |
| CVE-2024-52046 | The ObjectSerializationDecoder in Apache MINA uses Java's native deserialization protocol to process incoming serialized data but lacks the necessary security checks and defenses. This vulnerability allows attackers to exploit the deserialization process by sending specially crafted malicious serialized data, potentially leading to remote code execution (RCE) attacks. This issue affects MINA core versions 2.0.X, 2.1.X and 2.2.X, and will be fixed by the releases 2.0.27, 2.1.10 and 2.2.4. It's also important to note that an application using MINA core library will only be affected if the <code>IoBuffer#getObject()</code> method is called, and this specific method is potentially called when adding a <code>ProtocolCodecFilter</code> instance using the <code>ObjectSerializationCodecFactory</code> class in the filter chain. If your application is specifically using those classes, you have to upgrade to the latest version of MINA core library. Upgrading will not be enough: you also need to explicitly allow the classes the decoder will accept in the <code>ObjectSerializationDecoder</code> instance, using one of the three new methods: <code>/** * Accept class names where the supplied ClassNameMatcher matches for * deserialization, unless they are otherwise rejected. * * @param classNameMatcher the matcher to use */ public void accept(ClassNameMatcher classNameMatcher) /** * Accept class names that match the supplied pattern for * deserialization, unless they are otherwise rejected. * * @param pattern standard Java regexp */ public void accept(Pattern pattern) /** * Accept the wildcard specified classes for deserialization, * unless they are otherwise rejected. * * @param patterns Wildcard file name patterns as defined by * {@link org.apache.commons.io.FilenameUtils#wildcardMatch(String, String) FilenameUtils.wildcardMatch} */ public void accept(String... patterns)</code> By default, the decoder will reject <code>*all*</code> classes that will be present in the incoming data. Note: The <code>FtpServer</code> , <code>SSHd</code> and <code>Vysper</code> sub-project are not affected by this issue. |
| CVE-2024-52316 | Unchecked Error Condition vulnerability in Apache Tomcat. If Tomcat is configured to use a custom Jakarta Authentication (formerly JASPIC) <code>ServerAuthContext</code> component which may throw an exception during the authentication process without explicitly setting an HTTP status to indicate failure, the authentication may not fail, allowing the user to bypass the authentication process. There are no known Jakarta Authentication components that behave in this way. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.0-M26, from 10.1.0-M1 through 10.1.30, from 9.0.0-M1 through 9.0.95. Users are recommended to upgrade to version 11.0.0, 10.1.31 or 9.0.96, which fix the issue. |
| CVE-2024-53990 | The <code>AsyncHttpClient</code> (AHC) library allows Java applications to easily execute HTTP requests and asynchronously process HTTP responses. When making any HTTP request, the automatically enabled and self-managed <code>CookieStore</code> (aka cookie jar) will silently replace explicitly defined Cookies with any that have the same name from the cookie jar. For services that operate with multiple users, this can result in one user's Cookie being used for another user's requests. |
| CVE-2025-24813 | Path Equivalence: 'file.Name' (Internal Dot) leading to Remote Code Execution and/or Information disclosure and/or malicious content added to uploaded files via write enabled Default Servlet in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.2, from 10.1.0-M1 through 10.1.34, from 9.0.0.M1 through 9.0.98. The following versions were EOL at the time the CVE was created but are known to be affected: 8.5.0 through 8.5.100. Other, older, EOL versions may also be affected. If all of the following were true, a malicious user was able to view security sensitive files and/or inject content into those files: - writes enabled for the default servlet (disabled by default) - support for partial PUT (enabled by default) - a target URL for security sensitive uploads that was a sub-directory of a target URL for public uploads - attacker knowledge of the names of security sensitive files being uploaded - the security sensitive files also being uploaded via partial PUT If all of the following were true, a malicious user was able to perform remote code execution: - writes enabled for the default servlet (disabled by default) - support for partial PUT (enabled by default) - application was using Tomcat's file based session persistence with the default storage location - application included a library that may be leveraged in a deserialization attack Users are recommended to upgrade to version 11.0.3, 10.1.35 or 9.0.99, which fixes the issue. |
| CVE-2025-31651 | Improper Neutralization of Escape, Meta, or Control Sequences vulnerability in Apache Tomcat. For a subset of unlikely rewrite rule configurations, it was possible for a specially crafted request to bypass some rewrite rules. If those rewrite rules effectively enforced security constraints, those constraints could be bypassed. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.5, from 10.1.0-M1 through 10.1.39, from 9.0.0.M1 through 9.0.102. The following versions were EOL at the time the CVE was created but are known to be affected: 8.5.0 through 8.5.100. Other, older, EOL versions may also be affected. Users are recommended to upgrade to version [FIXED_VERSION], which fixes the issue. |

| CVE ID | Description |
|----------------------------------|---|
| CVE-2024-48910 | DOMPurify is a DOM-only, super-fast, uber-tolerant XSS sanitizer for HTML, MathML and SVG. DOMPurify was vulnerable to prototype pollution. This vulnerability is fixed in 2.4.2. |
| CVE-2024-45216 | Improper Authentication vulnerability in Apache Solr. Solr instances using the PKIAuthenticationPlugin, which is enabled by default when Solr Authentication is used, are vulnerable to Authentication bypass. A fake ending at the end of any Solr API URL path, will allow requests to skip Authentication while maintaining the API contract with the original URL Path. This fake ending looks like an unprotected API path, however it is stripped off internally after authentication but before API routing. This issue affects Apache Solr: from 5.3.0 before 8.11.4, from 9.0.0 before 9.7.0. Users are recommended to upgrade to version 9.7.0, or 8.11.4, which fix the issue. |
| CVE-2026-27727 | mchange-commons-java, a library that provides Java utilities, includes code that mirrors early implementations of JNDI functionality, including support for remote `factoryClassLocation` values, by which code can be downloaded and invoked within a running application. If an attacker can provoke an application to read a maliciously crafted `javax.naming.Reference` or serialized object, they can provoke the download and execution of malicious code. Implementations of this functionality within the JDK were disabled by default behind a System property that defaults to `false`, `com.sun.jndi.ldap.object.trustURLCodebase`. However, since mchange-commons-java includes an independent implementation of JNDI dereferencing, libraries (such as c3p0) that resolve references via that implementation could be provoked to download and execute malicious code even after the JDK was hardened. Mirroring the JDK patch, mchange-commons-java's JNDI functionality is gated by configuration parameters that default to restrictive values starting in version 0.4.0. No known workarounds are available. Versions prior to 0.4.0 should be avoided on application CLASSPATHs. |
| CVE-2025-66614 | Improper Input Validation vulnerability. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.14, from 10.1.0-M1 through 10.1.49, from 9.0.0-M1 through 9.0.112. The following versions were EOL at the time the CVE was created but are known to be affected: 8.5.0 through 8.5.100. Older EOL versions are not affected. Tomcat did not validate that the host name provided via the SNI extension was the same as the host name provided in the HTTP host header field. If Tomcat was configured with more than one virtual host and the TLS configuration for one of those hosts did not require client certificate authentication but another one did, it was possible for a client to bypass the client certificate authentication by sending different host names in the SNI extension and the HTTP host header field. The vulnerability only applies if client certificate authentication is only enforced at the Connector. It does not apply if client certificate authentication is enforced at the web application. Users are recommended to upgrade to version 11.0.15 or later, 10.1.50 or later or 9.0.113 or later, which fix the issue. |
| CVE-2025-24813 | Path Equivalence: `file.Name` (Internal Dot) leading to Remote Code Execution and/or Information disclosure and/or malicious content added to uploaded files via write enabled Default Servlet in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.2, from 10.1.0-M1 through 10.1.34, from 9.0.0-M1 through 9.0.98. The following versions were EOL at the time the CVE was created but are known to be affected: 8.5.0 through 8.5.100. Other, older, EOL versions may also be affected. If all of the following were true, a malicious user was able to view security sensitive files and/or inject content into those files: - writes enabled for the default servlet (disabled by default) - support for partial PUT (enabled by default) - a target URL for security sensitive uploads that was a sub-directory of a target URL for public uploads - attacker knowledge of the names of security sensitive files being uploaded - the security sensitive files also being uploaded via partial PUT If all of the following were true, a malicious user was able to perform remote code execution: - writes enabled for the default servlet (disabled by default) - support for partial PUT (enabled by default) - application was using Tomcat's file based session persistence with the default storage location - application included a library that may be leveraged in a deserialization attack Users are recommended to upgrade to version 11.0.3, 10.1.35 or 9.0.99, which fixes the issue. |
| CVE-2017-1000028 | Oracle, GlassFish Server Open Source Edition 4.1 is vulnerable to both authenticated and unauthenticated Directory Traversal vulnerability, that can be exploited by issuing a specially crafted HTTP GET request. |
| CVE-2019-18218 | <code>cdf_read_property_info</code> in <code>cdf.c</code> in file through 5.37 does not restrict the number of <code>CDF_VECTOR</code> elements, which allows a heap-based buffer overflow (4-byte out-of-bounds write). |
| CVE-2021-0341 | In <code>verifyHostName</code> of <code>OkHostnameVerifier.java</code> , there is a possible way to accept a certificate for the wrong domain due to improperly used crypto. This could lead to remote information disclosure with no additional execution privileges needed. User interaction is not needed for exploitation. Product: Android Versions: Android-8.1 Android-9 Android-10 Android-11 Android ID: A-171980069 |
| CVE-2021-44878 | If an OpenID Connect provider supports the "none" algorithm (i.e., tokens with no signature), <code>pac4j v5.3.0</code> (and prior) does not refuse it without an explicit configuration on its side or for the "idtoken" response type which is not secure and violates the OpenID Core Specification. The "none" algorithm does not require any signature verification when validating the ID tokens, which allows the attacker to bypass the token validation by injecting a malformed ID token using "none" as the value of "alg" key in the header with an empty signature value. |
| CVE-2022-25844 | The package <code>angular</code> after 1.7.0 are vulnerable to Regular Expression Denial of Service (ReDoS) by providing a custom locale rule that makes it possible to assign the parameter in <code>posPre: ''.repeat()</code> of <code>NUMBER_FORMATS.PATTERNS[1].posPre</code> with a very high value. Note: 1) This package has been deprecated and is no longer maintained. 2) The vulnerable versions are 1.7.0 and higher. |

| CVE ID | Description |
|--------------------------------|---|
| CVE-2022-3171 | A parsing issue with binary data in protobuf-java core and lite versions prior to 3.21.7, 3.20.3, 3.19.6 and 3.16.3 can lead to a denial of service attack. Inputs containing multiple instances of non-repeated embedded messages with repeated or unknown fields causes objects to be converted back-n-forth between mutable and immutable forms, resulting in potentially long garbage collection pauses. We recommend updating to the versions mentioned above. |
| CVE-2022-3509 | A parsing issue similar to CVE-2022-3171, but with textformat in protobuf-java core and lite versions prior to 3.21.7, 3.20.3, 3.19.6 and 3.16.3 can lead to a denial of service attack. Inputs containing multiple instances of non-repeated embedded messages with repeated or unknown fields causes objects to be converted back-n-forth between mutable and immutable forms, resulting in potentially long garbage collection pauses. We recommend updating to the versions mentioned above. |
| CVE-2022-3510 | A parsing issue similar to CVE-2022-3171, but with Message-Type Extensions in protobuf-java core and lite versions prior to 3.21.7, 3.20.3, 3.19.6 and 3.16.3 can lead to a denial of service attack. Inputs containing multiple instances of non-repeated embedded messages with repeated or unknown fields causes objects to be converted back-n-forth between mutable and immutable forms, resulting in potentially long garbage collection pauses. We recommend updating to the versions mentioned above. |
| CVE-2023-3635 | GzipSource does not handle an exception that might be raised when parsing a malformed gzip buffer. This may lead to denial of service of the Okio client when handling a crafted GZIP archive, by using the GzipSource class. |
| CVE-2023-39410 | When deserializing untrusted or corrupted data, it is possible for a reader to consume memory beyond the allowed constraints and thus lead to out of memory on the system. This issue affects Java applications using Apache Avro Java SDK up to and including 1.11.2. Users should update to apache-avro version 1.11.3 which addresses this issue. |
| CVE-2023-4586 | A vulnerability was found in the Hot Rod client. This security issue occurs as the Hot Rod client does not enable hostname validation when using TLS, possibly resulting in a man-in-the-middle (MITM) attack. |
| CVE-2023-7272 | In Eclipse Parsson before 1.0.4 and 1.1.3, a document with a large depth of nested objects can allow an attacker to cause a Java stack overflow exception and denial of service. Eclipse Parsson allows processing (e.g. parse, generate, transform and query) JSON documents. |
| CVE-2024-13009 | In Eclipse Jetty versions 9.4.0 to 9.4.56 a buffer can be incorrectly released when confronted with a gzip error when inflating a request body. This can result in corrupted and/or inadvertent sharing of data between requests. |
| CVE-2024-21490 | This affects versions of the package angular from 1.3.0. A regular expression used to split the value of the ng-srcset directive is vulnerable to super-linear runtime due to backtracking. With large carefully-crafted input, this can result in catastrophic backtracking and cause a denial of service. Note: This package is EOL and will not receive any updates to address this issue. Users should migrate to [@angular/core](https://www.npmjs.com/package/@angular/core) . |
| CVE-2024-25638 | dnsjava is an implementation of DNS in Java. Records in DNS replies are not checked for their relevance to the query, allowing an attacker to respond with RRs from different zones. This vulnerability is fixed in 3.6.0. |
| CVE-2024-29131 | Out-of-bounds Write vulnerability in Apache Commons Configuration. This issue affects Apache Commons Configuration: from 2.0 before 2.10.1. Users are recommended to upgrade to version 2.10.1, which fixes the issue. |
| CVE-2024-29857 | An issue was discovered in ECCurve.java and ECCurve.cs in Bouncy Castle Java (BC Java) before 1.78, BC Java LTS before 2.73.6, BC-FJA before 1.0.2.5, and BC C#.Net before 2.3.1. Importing an EC certificate with crafted F2m parameters can lead to excessive CPU consumption during the evaluation of the curve parameters. |
| CVE-2024-34447 | An issue was discovered in the Bouncy Castle Crypto Package For Java before BC TLS Java 1.0.19 (ships with BC Java 1.78, BC Java (LTS) 2.73.6) and before BC FIPS TLS Java 1.0.19. When endpoint identification is enabled in the BCJSSE and an SSL socket is created without an explicit hostname (as happens with <code>HttpsURLConnection</code>), hostname verification could be performed against a DNS-resolved IP address in some situations, opening up a possibility of DNS poisoning. |
| CVE-2024-47561 | Schema parsing in the Java SDK of Apache Avro 1.11.3 and previous versions allows bad actors to execute arbitrary code. Users are recommended to upgrade to version 1.11.4 or 1.12.0, which fix this issue. |
| CVE-2024-57699 | A security issue was found in Netplex Json-smart 2.5.0 through 2.5.1. When loading a specially crafted JSON input, containing a large number of '{', a stack exhaustion can be triggered, which could allow an attacker to cause a Denial of Service (DoS). This issue exists because of an incomplete fix for CVE-2023-1370. |
| CVE-2024-7254 | Any project that parses untrusted Protocol Buffers data containing an arbitrary number of nested groups / series of SGROUP tags can be corrupted by exceeding the stack limit i.e. StackOverflow. Parsing nested groups as unknown fields with <code>DiscardUnknownFieldsParser</code> or Java Protobuf Lite parser, or against Protobuf map fields, creates unbounded recursions that can be abused by an attacker. |
| CVE-2025-1948 | In Eclipse Jetty versions 12.0.0 to 12.0.16 included, an HTTP/2 client can specify a very large value for the HTTP/2 settings parameter <code>SETTINGS_MAX_HEADER_LIST_SIZE</code> . The Jetty HTTP/2 server does not perform validation on this setting, and tries to allocate a <code>ByteBuffer</code> of the specified capacity to encode HTTP responses, likely resulting in <code>OutOfMemoryError</code> being thrown, or even the JVM process exiting. |

| CVE ID | Description |
|--------------------------------|--|
| CVE-2025-22228 | BCryptPasswordEncoder.matches(CharSequence,String) will incorrectly return true for passwords larger than 72 characters as long as the first 72 characters are the same. |
| CVE-2025-24970 | Netty, an asynchronous, event-driven network application framework, has a vulnerability starting in version 4.1.91.Final and prior to version 4.1.118.Final. When a special crafted packet is received via SslHandler it doesn't correctly handle validation of such a packet in all cases which can lead to a native crash. Version 4.1.118.Final contains a patch. As workaround its possible to either disable the usage of the native SslEngine or change the code manually. |
| CVE-2025-31650 | Improper Input Validation vulnerability in Apache Tomcat. Incorrect error handling for some invalid HTTP priority headers resulted in incomplete clean-up of the failed request which created a memory leak. A large number of such requests could trigger an OutOfMemoryException resulting in a denial of service. This issue affects Apache Tomcat: from 9.0.76 through 9.0.102, from 10.1.10 through 10.1.39, from 11.0.0-M2 through 11.0.5. The following versions were EOL at the time the CVE was created but are known to be affected: 8.5.90 though 8.5.100. Users are recommended to upgrade to version 9.0.104, 10.1.40 or 11.0.6 which fix the issue. |
| CVE-2025-46762 | Schema parsing in the parquet-avro module of Apache Parquet 1.15.0 and previous versions allows bad actors to execute arbitrary code. While 1.15.1 introduced a fix to restrict untrusted packages, the default setting of trusted packages still allows malicious classes from these packages to be executed. The exploit is only applicable if the client code of parquet-avro uses the "specific" or the "reflect" models deliberately for reading Parquet files. ("generic" model is not impacted) Users are recommended to upgrade to 1.15.2 or set the system property "org.apache.parquet.avro.SERIALIZABLE_PACKAGES" to an empty string on 1.15.1. Both are sufficient to fix the issue. |
| CVE-2025-67721 | Aircompressor is a library with ports of the Snappy, LZ0, LZ4, and Zstandard compression algorithms to Java. In versions 3.3 and below, incorrect handling of malformed data in Java-based decompressor implementations for Snappy and LZ4 allow remote attackers to read previous buffer contents via crafted compressed input. With certain crafted compressed inputs, elements from the output buffer can end up in the uncompressed output, potentially leaking sensitive data. This is relevant for applications that reuse the same output buffer to uncompress multiple inputs. This can be the case of a web server that allocates a fix-sized buffer for performance purposes. There is similar vulnerability in GHSA-cmp6-m4wj-q63q. This issue is fixed in version 3.4. |
| CVE-2025-55163 | Netty is an asynchronous, event-driven network application framework. Prior to versions 4.1.124.Final and 4.2.4.Final, Netty is vulnerable to MadeYouReset DDoS. This is a logical vulnerability in the HTTP/2 protocol, that uses malformed HTTP/2 control frames in order to break the max concurrent streams limit - which results in resource exhaustion and distributed denial of service. This issue has been patched in versions 4.1.124.Final and 4.2.4.Final. |
| CVE-2025-58057 | Netty is an asynchronous event-driven network application framework for rapid development of maintainable high performance protocol servers & clients. In netty-codec-compression versions 4.1.124.Final and below, and netty-codec versions 4.2.4.Final and below, when supplied with specially crafted input, BrotliDecoder and certain other decompression decoders will allocate a large number of reachable byte buffers, which can lead to denial of service. BrotliDecoder.decompress has no limit in how often it calls pull, decompressing data 64K bytes at a time. The buffers are saved in the output list, and remain reachable until OOM is hit. This is fixed in versions 4.1.125.Final of netty-codec and 4.2.5.Final of netty-codec-compression. |
| CVE-2025-58056 | Netty is an asynchronous event-driven network application framework for development of maintainable high performance protocol servers and clients. In versions 4.1.124.Final, and 4.2.0.Alpha3 through 4.2.4.Final, Netty incorrectly accepts standalone newline characters (LF) as a chunk-size line terminator, regardless of a preceding carriage return (CR), instead of requiring CRLF per HTTP/1.1 standards. When combined with reverse proxies that parse LF differently (treating it as part of the chunk extension), attackers can craft requests that the proxy sees as one request but Netty processes as two, enabling request smuggling attacks. This is fixed in versions 4.1.125.Final and 4.2.5.Final. |
| CVE-2025-48734 | Improper Access Control vulnerability in Apache Commons. A special BeanIntrospector class was added in version 1.9.2. This can be used to stop attackers from using the declared class property of Java enum objects to get access to the classloader. However this protection was not enabled by default. PropertyUtilsBean (and consequently BeanUtilsBean) now disallows declared class level property access by default. Releases 1.11.0 and 2.0.0-M2 address a potential security issue when accessing enum properties in an uncontrolled way. If an application using Commons BeanUtils passes property paths from an external source directly to the getProperty() method of PropertyUtilsBean, an attacker can access the enum's class loader via the "declaredClass" property available on all Java "enum" objects. Accessing the enum's "declaredClass" allows remote attackers to access the ClassLoader and execute arbitrary code. The same issue exists with PropertyUtilsBean.getNestedProperty(). Starting in versions 1.11.0 and 2.0.0-M2 a special BeanIntrospector suppresses the "declaredClass" property. Note that this new BeanIntrospector is enabled by default, but you can disable it to regain the old behavior; see section 2.5 of the user's guide and the unit tests. This issue affects Apache Commons BeanUtils 1.x before 1.11.0, and 2.x before 2.0.0-M2. Users of the artifact commons-beanutils:commons-beanutils 1.x are recommended to upgrade to version 1.11.0, which fixes the issue. Users of the artifact org.apache.commons:commons-beanutils2 2.x are recommended to upgrade to version 2.0.0-M2, which fixes the issue. |

| CVE ID | Description |
|----------------|---|
| CVE-2026-24308 | Improper handling of configuration values in ZKConfig in Apache ZooKeeper 3.8.5 and 3.9.4 on all platforms allows an attacker to expose sensitive information stored in client configuration in the client's logfile. Configuration values are exposed at INFO level logging rendering potential production systems affected by the issue. Users are recommended to upgrade to version 3.8.6 or 3.9.5 which fixes this issue. |
| CVE-2026-24281 | Hostname verification in Apache ZooKeeper ZKTrustManager falls back to reverse DNS (PTR) when IP SAN validation fails, allowing attackers who control or spoof PTR records to impersonate ZooKeeper servers or clients with a valid certificate for the PTR name. It's important to note that attacker must present a certificate which is trusted by ZKTrustManager which makes the attack vector harder to exploit. Users are recommended to upgrade to version 3.8.6 or 3.9.5, which fixes this issue by introducing a new configuration option to disable reverse DNS lookup in client and quorum protocols. |
| CVE-2023-1428 | There exists a vulnerability causing an abort() to be called in gRPC. The following headers cause gRPC's C++ implementation to abort() when called via http2: te: x (x != trailers) :scheme: x (x != http, https) grpclb_client_stats: x (x == anything) On top of sending one of those headers, a later header must be sent that gets the total header size past 8KB. We recommend upgrading past git commit 2485fa94bd8a723e5c977d55a3ce10b301b437f8 or v1.53 and above. |
| CVE-2026-22022 | Deployments of Apache Solr 5.3.0 through 9.10.0 that rely on Solr's "Rule Based Authorization Plugin" are vulnerable to allowing unauthorized access to certain Solr APIs, due to insufficiently strict input validation in those components. Only deployments that meet all of the following criteria are impacted by this vulnerability: * Use of Solr's "RuleBasedAuthorizationPlugin" * A RuleBasedAuthorizationPlugin config (see security.json) that specifies multiple "roles" * A RuleBasedAuthorizationPlugin permission list (see security.json) that uses one or more of the following pre-defined permission rules: "config-read", "config-edit", "schema-read", "metrics-read", or "security-read". * A RuleBasedAuthorizationPlugin permission list that doesn't define the "all" pre-defined permission * A networking setup that allows clients to make unfiltered network requests to Solr. (i.e. user-submitted HTTP/HTTPS requests reach Solr as-is, unmodified or restricted by any intervening proxy or gateway) Users can mitigate this vulnerability by ensuring that their RuleBasedAuthorizationPlugin configuration specifies the "all" pre-defined permission and associates the permission with an "admin" or other privileged role. Users can also upgrade to a Solr version outside of the impacted range, such as the recently released Solr 9.10.1. |
| CVE-2024-45217 | Insecure Default Initialization of Resource vulnerability in Apache Solr. New ConfigSets that are created via a Restore command, which copy a configSet from the backup and give it a new name, are created without setting the "trusted" metadata. ConfigSets that do not contain the flag are trusted implicitly if the metadata is missing, therefore this leads to "trusted" ConfigSets that may not have been created with an Authenticated request. "trusted" ConfigSets are able to load custom code into classloaders, therefore the flag is supposed to only be set when the request that uploads the ConfigSet is Authenticated & Authorized. This issue affects Apache Solr: from 6.6.0 before 8.11.4, from 9.0.0 before 9.7.0. This issue does not affect Solr instances that are secured via Authentication/Authorization. Users are primarily recommended to use Authentication and Authorization when running Solr. However, upgrading to version 9.7.0, or 8.11.4 will mitigate this issue otherwise. |
| CVE-2026-22444 | The "create core" API of Apache Solr 8.6 through 9.10.0 lacks sufficient input validation on some API parameters, which can cause Solr to check the existence of and attempt to read file-system paths that should be disallowed by Solr's "allowPaths" security setting https://solr.apache.org/guide/solr/latest/configuration-guide/configuring-solr-xml.html#the-solr-element . These read-only accesses can allow users to create cores using unexpected configs if any are accessible via the filesystem. On Windows systems configured to allow UNC paths this can additionally cause disclosure of NTLM "user" hashes. Solr deployments are subject to this vulnerability if they meet the following criteria: * Solr is running in its "standalone" mode. * Solr's "allowPath" setting is being used to restrict file access to certain directories. * Solr's "create core" API is exposed and accessible to untrusted users. This can happen if Solr's RuleBasedAuthorizationPlugin https://solr.apache.org/guide/solr/latest/deployment-guide/rule-based-authorization-plugin.html is disabled, or if it is enabled but the "core-admin-edit" predefined permission (or an equivalent custom permission) is given to low-trust (i.e. non-admin) user roles. Users can mitigate this by enabling Solr's RuleBasedAuthorizationPlugin (if disabled) and configuring a permission-list that prevents untrusted users from creating new Solr cores. Users should also upgrade to Apache Solr 9.10.1 or greater, which contain fixes for this issue. |
| CVE-2025-5222 | A stack buffer overflow was found in International components for unicode (ICU). While running the genrb binary, the 'subtag' struct overflowed at the SRBRoot::addTag function. This issue may lead to memory corruption and local arbitrary code execution. |
| CVE-2025-5115 | In Eclipse Jetty, versions <=9.4.57, <=10.0.25, <=11.0.25, <=12.0.21, <=12.1.0.alpha2, an HTTP/2 client may trigger the server to send RST_STREAM frames, for example by sending frames that are malformed or that should not be sent in a particular stream state, therefore forcing the server to consume resources such as CPU and memory. For example, a client can open a stream and then send WINDOW_UPDATE frames with window size increment of 0, which is illegal. Per specification https://www.rfc-editor.org/rfc/rfc9113.html#name-window_update , the server should send a RST_STREAM frame. The client can now open another stream and send another bad WINDOW_UPDATE, therefore causing the server to consume more resources than necessary, as this case does not exceed the max number of concurrent streams, yet the client is able to create an enormous amount of streams in a short period of time. The attack can be performed with other conditions (for example, a DATA frame for a closed stream) that cause the server to send a RST_STREAM frame. Links: * https://github.com/jetty/jetty.project/security/advisories/GHSA-mmxm-8w33-wc4h |

| CVE ID | Description |
|--------------------------------|---|
| CVE-2019-10768 | In AngularJS before 1.7.9 the function <code>merge()</code> could be tricked into adding or modifying properties of <code>Object.prototype</code> using a <code>__proto__</code> payload. |
| CVE-2022-2048 | In Eclipse Jetty HTTP/2 server implementation, when encountering an invalid HTTP/2 request, the error handling has a bug that can wind up not properly cleaning up the active connections and associated resources. This can lead to a Denial of Service scenario where there are no enough resources left to process good requests. |
| CVE-2023-36478 | Eclipse Jetty provides a web server and servlet container. In versions 11.0.0 through 11.0.15, 10.0.0 through 10.0.15, and 9.0.0 through 9.4.52, an integer overflow in <code>MetaDataBuilder.checkSize</code> allows for HTTP/2 HPACK header values to exceed their size limit. <code>MetaDataBuilder.java</code> determines if a header name or value exceeds the size limit, and throws an exception if the limit is exceeded. However, when length is very large and huffman is true, the multiplication by 4 in line 295 will overflow, and length will become negative. <code>_(size+length)</code> will now be negative, and the check on line 296 will not be triggered. Furthermore, <code>MetaDataBuilder.checkSize</code> allows for user-entered HPACK header value sizes to be negative, potentially leading to a very large buffer allocation later on when the user-entered size is multiplied by 2. This means that if a user provides a negative length value (or, more precisely, a length value which, when multiplied by the 4/3 fudge factor, is negative), and this length value is a very large positive number when multiplied by 2, then the user can cause a very large buffer to be allocated on the server. Users of HTTP/2 can be impacted by a remote denial of service attack. The issue has been fixed in versions 11.0.16, 10.0.16, and 9.4.53. There are no known workarounds. |
| CVE-2023-44487 | The HTTP/2 protocol allows a denial of service (server resource consumption) because request cancellation can reset many streams quickly, as exploited in the wild in August through October 2023. |
| CVE-2024-22201 | Jetty is a Java based web server and servlet engine. An HTTP/2 SSL connection that is established and TCP congested will be leaked when it times out. An attacker can cause many connections to end up in this state, and the server may run out of file descriptors, eventually causing the server to stop accepting new connections from valid clients. The vulnerability is patched in 9.4.54, 10.0.20, 11.0.20, and 12.0.6. |
| CVE-2024-9823 | There exists a security vulnerability in Jetty's <code>DosFilter</code> which can be exploited by unauthorized users to cause remote denial-of-service (DoS) attack on the server using <code>DosFilter</code> . By repeatedly sending crafted requests, attackers can trigger <code>OutOfMemory</code> errors and exhaust the server's memory finally. |
| CVE-2023-50386 | Improper Control of Dynamically-Managed Code Resources, Unrestricted Upload of File with Dangerous Type, Inclusion of Functionality from Untrusted Control Sphere vulnerability in Apache Solr. This issue affects Apache Solr: from 6.0.0 through 8.11.2, from 9.0.0 before 9.4.1. In the affected versions, Solr ConfigSets accepted Java jar and class files to be uploaded through the ConfigSets API. When backing up Solr Collections, these configSet files would be saved to disk when using the <code>LocalFileSystemRepository</code> (the default for backups). If the backup was saved to a directory that Solr uses in its <code>ClassPath/ClassLoaders</code> , then the jar and class files would be available to use with any ConfigSet, trusted or untrusted. When Solr is run in a secure way (Authorization enabled), as is strongly suggested, this vulnerability is limited to extending the Backup permissions with the ability to add libraries. Users are recommended to upgrade to version 8.11.3 or 9.4.1, which fix the issue. In these versions, the following protections have been added: * Users are no longer able to upload files to a configSet that could be executed via a Java ClassLoader. * The Backup API restricts saving backups to directories that are used in the ClassLoader. |
| CVE-2023-50291 | Insufficiently Protected Credentials vulnerability in Apache Solr. This issue affects Apache Solr: from 6.0.0 through 8.11.2, from 9.0.0 before 9.3.0. One of the two endpoints that publishes the Solr process' Java system properties, <code>/admin/info/properties</code> , was only setup to hide system properties that had "password" contained in the name. There are a number of sensitive system properties, such as "basicauth" and "aws.secretKey" do not contain "password", thus their values were published via the <code>/admin/info/properties</code> endpoint. This endpoint populates the list of System Properties on the home screen of the Solr Admin page, making the exposed credentials visible in the UI. This <code>/admin/info/properties</code> endpoint is protected under the "config-read" permission. Therefore, Solr Clouds with Authorization enabled will only be vulnerable through logged-in users that have the "config-read" permission. Users are recommended to upgrade to version 9.3.0 or 8.11.3, which fixes the issue. A single option now controls hiding Java system property for all endpoints, <code>"-Dsolr.hiddenSysProps"</code> . By default all known sensitive properties are hidden (including <code>"-Dbasicauth"</code>), as well as any property with a name containing "secret" or "password". Users who cannot upgrade can also use the following Java system property to fix the issue: <code>"-Dsolr.redaction.system.pattern=.*(password secret basicauth).*"</code> |
| CVE-2023-50292 | Incorrect Permission Assignment for Critical Resource, Improper Control of Dynamically-Managed Code Resources vulnerability in Apache Solr. This issue affects Apache Solr: from 8.10.0 through 8.11.2, from 9.0.0 before 9.3.0. The Schema Designer was introduced to allow users to more easily configure and test new Schemas and configSets. However, when the feature was created, the "trust" (authentication) of these configSets was not considered. External library loading is only available to configSets that are "trusted" (created by authenticated users), thus non-authenticated users are unable to perform Remote Code Execution. Since the Schema Designer loaded configSets without taking their "trust" into account, configSets that were created by unauthenticated users were allowed to load external libraries when used in the Schema Designer. Users are recommended to upgrade to version 9.3.0, which fixes the issue. |

| CVE ID | Description |
|----------------|--|
| CVE-2023-50298 | Exposure of Sensitive Information to an Unauthorized Actor vulnerability in Apache Solr. This issue affects Apache Solr: from 6.0.0 through 8.11.2, from 9.0.0 before 9.4.1. Solr Streaming Expressions allows users to extract data from other Solr Clouds, using a "zkHost" parameter. When original SolrCloud is setup to use ZooKeeper credentials and ACLs, they will be sent to whatever "zkHost" the user provides. An attacker could setup a server to mock ZooKeeper, that accepts ZooKeeper requests with credentials and ACLs and extracts the sensitive information, then send a streaming expression using the mock server's address in "zkHost". Streaming Expressions are exposed via the "/streaming" handler, with "read" permissions. Users are recommended to upgrade to version 8.11.3 or 9.4.1, which fix the issue. From these versions on, only zkHost values that have the same server address (regardless of chroot), will use the given ZooKeeper credentials and ACLs when connecting. |
| CVE-2025-66524 | Apache NiFi 1.20.0 through 2.6.0 include the GetAsanaObject Processor, which requires integration with a configurable Distribute Map Cache Client Service for storing and retrieving state information. The GetAsanaObject Processor used generic Java Object serialization and deserialization without filtering. Unfiltered Java object deserialization does not provide protection against crafted state information stored in the cache server configured for GetAsanaObject. Exploitation requires an Apache NiFi system running with the GetAsanaObject Processor, and direct access to the configured cache server. Upgrading to Apache NiFi 2.7.0 is the recommended mitigation, which replaces Java Object serialization with JSON serialization. Removing the GetAsanaObject Processor located in the nifi-asana-processors-nar bundle also prevents exploitation. |
| CVE-2020-14734 | Vulnerability in the Oracle Text component of Oracle Database Server. Supported versions that are affected are 11.2.0.4, 12.1.0.2, 12.2.0.1, 18c and 19c. Difficult to exploit vulnerability allows unauthenticated attacker with network access via Oracle Net to compromise Oracle Text. Successful attacks of this vulnerability can result in takeover of Oracle Text. CVSS 3.1 Base Score 8.1 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:N/UI:N/S:U/C:H/I:H/A:H). |
| CVE-2017-15288 | The compilation daemon in Scala before 2.10.7, 2.11.x before 2.11.12, and 2.12.x before 2.12.4 uses weak permissions for private files in /tmp/scala-devel/\${USER:shared}/scalac-compile-server-port, which allows local users to write to arbitrary class files and consequently gain privileges. |
| CVE-2023-45142 | OpenTelemetry-Go Contrib is a collection of third-party packages for OpenTelemetry-Go. A handler wrapper out of the box adds labels `http.user_agent` and `http.method` that have unbound cardinality. It leads to the server's potential memory exhaustion when many malicious requests are sent to it. HTTP header User-Agent or HTTP method for requests can be easily set by an attacker to be random and long. The library internally uses `httpconv.ServerRequest` that records every value for HTTP `method` and `User-Agent`. In order to be affected, a program has to use the `otelhttp.NewHandler` wrapper and not filter any unknown HTTP methods or User agents on the level of CDN, LB, previous middleware, etc. Version 0.44.0 fixed this issue when the values collected for attribute `http.request.method` were changed to be restricted to a set of well-known values and other high cardinality attributes were removed. As a workaround to stop being affected, `otelhttp.WithFilter()` can be used, but it requires manual careful configuration to not log certain requests entirely. For convenience and safe usage of this library, it should by default mark with the label `unknown` non-standard HTTP methods and User agents to show that such requests were made but do not increase cardinality. In case someone wants to stay with the current behavior, library API should allow to enable it. |
| CVE-2023-47108 | OpenTelemetry-Go Contrib is a collection of third-party packages for OpenTelemetry-Go. Starting in version 0.37.0 and prior to version 0.46.0, the grpc Unary Server Interceptor out of the box adds labels `net.peer.sock.addr` and `net.peer.sock.port` that have unbound cardinality. It leads to the server's potential memory exhaustion when many malicious requests are sent. An attacker can easily flood the peer address and port for requests. Version 0.46.0 contains a fix for this issue. As a workaround to stop being affected, a view removing the attributes can be used. The other possibility is to disable grpc metrics instrumentation by passing `otelgrpc.WithMeterProvider` option with `noop.NewMeterProvider`. |
| CVE-2025-54920 | This issue affects Apache Spark: before 3.5.7 and 4.0.1. Users are recommended to upgrade to version 3.5.7 or 4.0.1 and above, which fixes the issue. Summary Apache Spark 3.5.4 and earlier versions contain a code execution vulnerability in the Spark History Web UI due to overly permissive Jackson deserialization of event log data. This allows an attacker with access to the Spark event logs directory to inject malicious JSON payloads that trigger deserialization of arbitrary classes, enabling command execution on the host running the Spark History Server. Details The vulnerability arises because the Spark History Server uses Jackson polymorphic deserialization with @JsonTypeInfo.Id.CLASS on SparkListenerEvent objects, allowing an attacker to specify arbitrary class names in the event JSON. This behavior permits instantiating unintended classes, such as org.apache.hive.jdbc.HiveConnection, which can perform network calls or other malicious actions during deserialization. The attacker can exploit this by injecting crafted JSON content into the Spark event log files, which the History Server then deserializes on startup or when loading event logs. For example, the attacker can force the History Server to open a JDBC connection to a remote attacker-controlled server, demonstrating remote command injection capability. Proof of Concept: 1. Run Spark with event logging enabled, writing to a writable directory (spark-logs). 2. Inject the following JSON at the beginning of an event log file: { "Event": "org.apache.hive.jdbc.HiveConnection", "uri": "jdbc:hive2://<IP>:<PORT>/", "info": { "hive.metastore.uris": "thrift://<IP>:<PORT>" } } 3. Start the Spark History Server with logs pointing to the modified directory. 4. The Spark History Server initiates a JDBC connection to the attacker's server, confirming the injection. Impact An attacker with write access to Spark event logs can execute arbitrary code on the server running the History Server, potentially compromising the entire system. |

| CVE ID | Description |
|--------------------------------|--|
| CVE-2025-54988 | Critical XXE in Apache Tika (tika-parser-pdf-module) in Apache Tika 1.13 through and including 3.2.1 on all platforms allows an attacker to carry out XML External Entity injection via a crafted XFA file inside of a PDF. An attacker may be able to read sensitive data or trigger malicious requests to internal resources or third-party servers. Note that the tika-parser-pdf-module is used as a dependency in several Tika packages including at least: tika-parsers-standard-modules, tika-parsers-standard-package, tika-app, tika-grpc and tika-server-standard. Users are recommended to upgrade to version 3.2.2, which fixes this issue. |
| CVE-2026-24734 | Improper Input Validation vulnerability in Apache Tomcat Native, Apache Tomcat. When using an OCSP responder, Tomcat Native (and Tomcat's FFM port of the Tomcat Native code) did not complete verification or freshness checks on the OCSP response which could allow certificate revocation to be bypassed. This issue affects Apache Tomcat Native: from 1.3.0 through 1.3.4, from 2.0.0 through 2.0.11; Apache Tomcat: from 11.0.0-M1 through 11.0.17, from 10.1.0-M7 through 10.1.51, from 9.0.83 through 9.0.114. The following versions were EOL at the time the CVE was created but are known to be affected: from 1.1.23 through 1.1.34, from 1.2.0 through 1.2.39. Older EOL versions are not affected. Apache Tomcat Native users are recommended to upgrade to versions 1.3.5 or later or 2.0.12 or later, which fix the issue. Apache Tomcat users are recommended to upgrade to versions 11.0.18 or later, 10.1.52 or later or 9.0.115 or later which fix the issue. |
| CVE-2026-25639 | Axios is a promise based HTTP client for the browser and Node.js. Prior to versions 0.30.3 and 1.13.5, the mergeConfig function in axios crashes with a TypeError when processing configuration objects containing <code>__proto__</code> as an own property. An attacker can trigger this by providing a malicious configuration object created via <code>JSON.parse()</code> , causing complete denial of service. This vulnerability is fixed in versions 0.30.3 and 1.13.5. |
| CVE-2025-11965 | In Eclipse Vert.x versions [4.0.0, 4.5.21] and [5.0.0, 5.0.4], a StaticHandler configuration for restricting access to hidden files fails to restrict access to hidden directories, allowing unauthorized users to retrieve files within them (e.g. <code>'git/config'</code>). |
| CVE-2017-18214 | The moment module before 2.19.3 for Node.js is prone to a regular expression denial of service via a crafted date string, a different vulnerability than CVE-2016-4055. |
| CVE-2022-24785 | Moment.js is a JavaScript date library for parsing, validating, manipulating, and formatting dates. A path traversal vulnerability impacts npm (server) users of Moment.js between versions 1.0.1 and 2.29.1, especially if a user-provided locale string is directly used to switch moment locale. This problem is patched in 2.29.2, and the patch can be applied to all affected versions. As a workaround, sanitize the user-provided locale name before passing it to Moment.js. |
| CVE-2022-1271 | An arbitrary file write vulnerability was found in GNU <code>gzip</code> 's <code>zgrep</code> utility. When <code>zgrep</code> is applied on the attacker's chosen file name (for example, a crafted file name), this can overwrite an attacker's content to an arbitrary attacker-selected file. This flaw occurs due to insufficient validation when processing filenames with two or more newlines where selected content and the target file names are embedded in crafted multi-line file names. This flaw allows a remote, low privileged attacker to force <code>zgrep</code> to write arbitrary files on the system. |
| CVE-2015-4035 | <code>scripts/xzgrep.in</code> in <code>xzgrep</code> 5.2.x before 5.2.0, before 5.0.0 does not properly process file names containing semicolons, which allows remote attackers to execute arbitrary code by having a user run <code>xzgrep</code> on a crafted file name. |

Cloudera products has a transitive dependencies on one of the AWS library that is reporting these vulnerabilities. However, Cloudera products do not use this library in any of the vulnerable code paths as noted in the CVE. As such Cloudera product remains not vulnerable to this. When fixes from AWS becomes available, our products will upgrade to the latest versions in upcoming releases. The following CVEs are not vulnerable:

| CVE ID | Description |
|--------------------------------|--|
| CVE-2025-55163 | Netty is an asynchronous, event-driven network application framework. Prior to versions 4.1.124.Final and 4.2.4.Final, Netty is vulnerable to MadeYouReset DDoS. This is a logical vulnerability in the HTTP/2 protocol, that uses malformed HTTP/2 control frames in order to break the max concurrent streams limit - which results in resource exhaustion and distributed denial of service. This issue has been patched in versions 4.1.124.Final and 4.2.4.Final. |
| CVE-2025-24970 | Netty, an asynchronous, event-driven network application framework, has a vulnerability starting in version 4.1.91.Final and prior to version 4.1.118.Final. When a special crafted packet is received via <code>SslHandler</code> it doesn't correctly handle validation of such a packet in all cases which can lead to a native crash. Version 4.1.118.Final contains a patch. As workaround its possible to either disable the usage of the native <code>SSL</code> Engine or change the code manually. |
| CVE-2025-58057 | Netty is an asynchronous event-driven network application framework for rapid development of maintainable high performance protocol servers & clients. In <code>netty-codec-compression</code> versions 4.1.124.Final and below, and <code>netty-codec</code> versions 4.2.4.Final and below, when supplied with specially crafted input, <code>BrotdiDecoder</code> and certain other decompression decoders will allocate a large number of reachable byte buffers, which can lead to denial of service. <code>BrotdiDecoder.decompress</code> has no limit in how often it calls <code>pull</code> , decompressing data 64K bytes at a time. The buffers are saved in the output list, and remain reachable until OOM is hit. This is fixed in versions 4.1.125.Final of <code>netty-codec</code> and 4.2.5.Final of <code>netty-codec-compression</code> . |

| CVE ID | Description |
|--------------------------------|---|
| CVE-2024-47535 | Netty is an asynchronous event-driven network application framework for rapid development of maintainable high performance protocol servers & clients. An unsafe reading of environment file could potentially cause a denial of service in Netty. When loaded on an Windows application, Netty attempts to load a file that does not exist. If an attacker creates such a large file, the Netty application crashes. This vulnerability is fixed in 4.1.115. |
| CVE-2025-25193 | Netty, an asynchronous, event-driven network application framework, has a vulnerability in versions up to and including 4.1.118.Final. An unsafe reading of environment file could potentially cause a denial of service in Netty. When loaded on an Windows application, Netty attempts to load a file that does not exist. If an attacker creates such a large file, the Netty application crash. A similar issue was previously reported as CVE-2024-47535. This issue was fixed, but the fix was incomplete in that null-bytes were not counted against the input limit. Commit d1fba62d3a47835d3fb35db8bd42ecc205a5386 contains an updated fix. |
| CVE-2025-58056 | Netty is an asynchronous event-driven network application framework for development of maintainable high performance protocol servers and clients. In versions 4.1.124.Final, and 4.2.0.Alpha3 through 4.2.4.Final, Netty incorrectly accepts standalone newline characters (LF) as a chunk-size line terminator, regardless of a preceding carriage return (CR), instead of requiring CRLF per HTTP/1.1 standards. When combined with reverse proxies that parse LF differently (treating it as part of the chunk extension), attackers can craft requests that the proxy sees as one request but Netty processes as two, enabling request smuggling attacks. This is fixed in versions 4.1.125.Final and 4.2.5.Final. |

API Modifications and Removals

All the APIs that have either been modified or removed for this release.

Components with API differences in Cloudera Runtime 7.3.2

The following is a list of all the APIs that have been modified or removed in this release of Cloudera Runtime. The APIs listed here are on a component by component basis and may include recommendations for how to manage the change.

API Compatibility changes in 7.3.2 for Hadoop

Removed or Modified APIs in Cloudera Runtime 7.3.2 for Hadoop and recommendations for how to handle them.

Apache Version of Hadoop in 7.3.1 was 3.1.1 and Apache Version of Hadoop in 7.3.2 is 3.4.2

Removed APIs in 7.3.2

The following Hadoop API endpoints have been removed from Cloudera Runtime 7.3.2. This change is a result of the fix for [CVE-2025-48924](#) and YARN's migration to common-lang 3.

[GetApplicationsRequest.getFinishRange](#)

Method Removed

Package Name

org.apache.hadoop.yarn.api.protocolrecords

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[YARN-10772](#)

Recommendation

Now returns Range<Long> instead of LongRange.

Recompilation Required?

Yes

[GetApplicationsRequest.getStartRange](#)

Method Removed

Package Name

org.apache.hadoop.yarn.api.protocolrecords

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[YARN-10772](#)

Recommendation

Now returns Range<Long> instead of LongRange.

Recompilation Required?

Yes

GetApplicationsRequest.newInstance

Method Removed

Package Name

org.apache.hadoop.yarn.api.protocolrecords

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[YARN-10772](#)

Recommendation

factory method parameter types changed

Recompilation Required?

Yes

GetApplicationsRequest.setFinishRange

Method Removed

Package Name

org.apache.hadoop.yarn.api.protocolrecords

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[YARN-10772](#)

Recommendation

parameter type changed: request.setFinishRange(new LongRange(startTime, endTime)) uses commons-lang. Now changed to request.setFinishRange(Range.between(startTime, endTime)); // commons-lang3

Recompilation Required?

Yes

GetApplicationsRequest.setStartRange

Method Removed

Package Name

org.apache.hadoop.yarn.api.protocolrecords

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[YARN-10772](#)

Recommendation

```
parameter type changed: request.setStartRange(new LongRange(startTime, endTime)) uses
commons-lang. Now changed to request.setStartRange(Range.between(startTime, endTime)); //
commons-lang3
```

Recompilation Required?

Yes

API Compatibility changes in 7.3.2 for HBase

Removed or Modified APIs in Cloudera Runtime 7.3.2 for HBase and recommendations for how to handle them.

Apache Version of HBase in 7.3.1.600 was 2.4.17 and Apache Version of HBase in 7.3.2 is 2.6.3.

Modified APIs in 7.3.2

The following APIs have been modified for HBase in Cloudera Runtime 7.3.2

DataBlockEncoding.createEncoder

Access level has been changed from protected to package-private.

Package Name

```
org.apache.hadoop.hbase.io.encoding
```

Effect

Recompilation of a client program may be terminated with the message:
`createEncoder (String)` has package-private access in
`org.apache.hadoop.hbase.io.encoding.DataBlockEncoding`.

Reason for change

[HBASE-27206](#)

Recommendation

Use `DataBlockEncoding.getEncoder` instead of `createEncoder`.

Recompilation Required?

YES

Bytes.toByteString

Return value type has been changed from `com.google.protobuf.ByteString` to
`org.apache.hbase.thirdparty.com.google.protobuf.ByteString`.

Package Name

```
org.apache.hadoop.hbase.util
```

Effect

This method has been removed because the return type is part of the method signature. A client program may be interrupted by `NoSuchMethodError` exception.

Reason for change

CDPD-33849

Recommendation

Method signature changed, return type changed to HBase thirdparty protobuf class. If a client application calls this method, it should import the HBase thirdparty protobuf class.

Recompilation Required?

YES

PBType.inputStreamFromByteRange

Return value type has been changed from `com.google.protobuf.CodedInputStream` to
`org.apache.hbase.thirdparty.com.google.protobuf.CodedInputStream`.

Package Name

```
org.apache.hadoop.hbase.types
```

Effect

This method has been removed because the return type is part of the method signature. A client program may be interrupted by `NoSuchMethodError` exception.

Reason for change

CDPD-33849

Recommendation

Method signature changed, return type changed to HBase thirdparty protobuf class. If a client application calls this method, it should import the HBase thirdparty protobuf class.

Recompilation Required?

YES

PBType.outputStreamFromByteRange

Return value type has been changed from `com.google.protobuf.CodedOutputStream` to `org.apache.hbase.thirdparty.com.google.protobuf.CodedOutputStream`.

Package Name

`org.apache.hadoop.hbase.types`

Effect

This method has been removed because the return type is part of the method signature. A client program may be interrupted by `NoSuchMethodError` exception.

Reason for change

CDPD-33849

Recommendation

Method signature changed, return type changed to HBase thirdparty protobuf class. If a client application calls this method, it should import the HBase thirdparty protobuf class.

Recompilation Required?

YES

Table.coprocessorService

Added `org.apache.hbase.thirdparty.com.google.protobuf.ServiceException` exception thrown.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: unreported exception `org.apache.hbase.thirdparty.com.google.protobuf.ServiceException` must be caught or declared to be thrown.

Reason for change

CDPD-33849

Recommendation

Method signature changed, thrown Exception type changed to HBase thirdparty protobuf class. If a client application calls this method, and handles this Exception makes sure to import it from the HBase thirdparty protobuf class(es).

Recompilation Required?

YES

ClusterMetrics

Abstract method `List getMasterTasks()` has been added to this interface.

Package Name

`org.apache.hadoop.hbase`

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `getMasterTasks()` in `org.apache.hadoop.hbase.ClusterMetrics`.

Reason for change

[HBASE-26730](#)

Recommendation

If a class extends `ClusterMetrics` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

ClusterMetrics

Abstract method List `getUnknownServerNames()` has been added to this interface.

Package Name

`org.apache.hadoop.hbase`

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `getUnknownServerNames()` in `org.apache.hadoop.hbase.ClusterMetrics`.

Reason for change

[HBASE-27104](#)

Recommendation

If a class extends `ClusterMetrics` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

ServerMetrics

Abstract method List `getTasks()` has been added to this interface.

Package Name

`org.apache.hadoop.hbase`

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `getTasks()` in `org.apache.hadoop.hbase.ServerMetrics`.

Reason for change

[HBASE-26730](#)

Recommendation

If a class extends `ServerMetrics` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Admin

Abstract method `BalanceResponse balance` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `balance (BalanceRequest)` in `org.apache.hadoop.hbase.client.Admin`.

Reason for change

[HBASE-26147](#)

Recommendation

If a class extends `Admin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Admin

Abstract method `Future<Void> flushAsync (TableName; List<byte[] >)` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `flushAsync (TableName; List<byte[] >)` in `org.apache.hadoop.hbase.client.Admin`.

Reason for change

[HBASE-26867](#)

Recommendation

If a class extends `Admin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Admin

Abstract method `void flushMasterStore ()` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `flushMasterStore ()` in `org.apache.hadoop.hbase.client.Admin`.

Reason for change

[HBASE-27028](#)

Recommendation

If a class extends `Admin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Admin

Abstract method `boolean isReplicationPeerEnabled (String)` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `isReplicationPeerEnabled (String)` in `org.apache.hadoop.hbase.client.Admin`.

Reason for change[HBASE-27448](#)**Recommendation**

If a class extends Admin in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Admin

Abstract method `boolean isReplicationPeerEnabled (String)` has been added to this interface.

Package Name`org.apache.hadoop.hbase.client`**Effect**

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `isReplicationPeerModificationEnabled ()` in `org.apache.hadoop.hbase.client.Admin`.

Reason for change[HBASE-27783](#)**Recommendation**

If a class extends Admin in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Admin

Abstract method `List<TableDescriptor> listTableDescriptorsByState (boolean)` has been added to this interface.

Package Name`org.apache.hadoop.hbase.client`**Effect**

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `listTableDescriptorsByState (boolean)` in `org.apache.hadoop.hbase.client.Admin`.

Reason for change[HBASE-27444](#)**Recommendation**

If a class extends Admin in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Admin

Abstract method `List<TableName> listTableNamesByState (boolean)` has been added to this interface.

Package Name`org.apache.hadoop.hbase.client`**Effect**

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `listTableNamesByState (boolean)` in `org.apache.hadoop.hbase.client.Admin`.

Reason for change

[HBASE-27444](#)**Recommendation**

If a class extends Admin in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Admin

Abstract method `Future<Void> modifyTableAsync (TableDescriptor; boolean)` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `modifyTableAsync (TableDescriptor; boolean)` in `org.apache.hadoop.hbase.client.Admin`.

Reason for change

[HBASE-25549](#)

Recommendation

If a class extends Admin in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Admin

Abstract method `boolean replicationPeerModificationSwitch (boolean; boolean)` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `replicationPeerModificationSwitch (boolean; boolean)` in `org.apache.hadoop.hbase.client.Admin`.

Reason for change

[HBASE-27783](#)

Recommendation

If a class extends Admin in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdmin

Abstract method `CompletableFuture<BalanceResponse> balance (BalanceRequest)` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `balance (BalanceRequest)` in `org.apache.hadoop.hbase.client.AsyncAdmin`.

Reason for change

[HBASE-26147](#)

Recommendation

If a class extends `AsyncAdmin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdmin

Abstract method `CompletableFuture<Void> flush (TableName; List<byte[]>)` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: a client class `C` is not abstract and does not override abstract method `flush (TableName; List<byte[]>)` in `org.apache.hadoop.hbase.client.AsyncAdmin`.

Reason for change

[HBASE-26867](#)

Recommendation

If a class extends `AsyncAdmin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdmin

Abstract method `CompletableFuture<Void> flushMasterStore ()` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: a client class `C` is not abstract and does not override abstract method `flushMasterStore ()` in `org.apache.hadoop.hbase.client.AsyncAdmin`.

Reason for change

[HBASE-27028](#)

Recommendation

If a class extends `AsyncAdmin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdmin

Abstract method `CompletableFuture<Boolean> isReplicationPeerEnabled (String)` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: a client class `C` is not abstract and does not override abstract method `isReplicationPeerEnabled (String)` in `org.apache.hadoop.hbase.client.AsyncAdmin`.

Reason for change

[HBASE-27448](#)

Recommendation

If a class extends `AsyncAdmin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdmin

Abstract method `CompletableFuture<Boolean> isReplicationPeerModificationEnabled ()` has been added to this interface.

Package Name

```
org.apache.hadoop.hbase.client
```

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `isReplicationPeerModificationEnabled ()` in `org.apache.hadoop.hbase.client.AsyncAdmin`.

Reason for change

[HBASE-27783](#)

Recommendation

If a class extends `AsyncAdmin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdmin

Abstract method `CompletableFuture<List<TableDescriptor>> listTableDescriptorsByState (boolean)` has been added to this interface.

Package Name

```
org.apache.hadoop.hbase.client
```

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `listTableDescriptorsByState (boolean)` in `org.apache.hadoop.hbase.client.AsyncAdmin`.

Reason for change

[HBASE-27444](#)

Recommendation

If a class extends `AsyncAdmin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdmin

Abstract method `CompletableFuture<List<TableName>> listTableNamesByState (boolean)` has been added to this interface.

Package Name

```
org.apache.hadoop.hbase.client
```

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `listTableNamesByState (boolean)` in `org.apache.hadoop.hbase.client.AsyncAdmin`.

Reason for change

[HBASE-27444](#)

Recommendation

If a class extends `AsyncAdmin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdmin

Abstract method `CompletableFuture<Void> modifyTable (TableDescriptor; boolean)` has been added to this interface.

Package Name`org.apache.hadoop.hbase.client`**Effect**

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `modifyTable (TableDescriptor; boolean)` in `org.apache.hadoop.hbase.client.AsyncAdmin`.

Reason for change[HBASE-25549](#)**Recommendation**

If a class extends `AsyncAdmin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdmin

Abstract method `CompletableFuture<Boolean> replicationPeerModificationSwitch (boolean; boolean)` has been added to this interface.

Package Name`org.apache.hadoop.hbase.client`**Effect**

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `replicationPeerModificationSwitch (boolean; boolean)` in `org.apache.hadoop.hbase.client.AsyncAdmin`.

Reason for change[HBASE-27783](#)**Recommendation**

If a class extends `AsyncAdmin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdminBuilder

Abstract method `AsyncAdminBuilder setRetryPauseForServerOverloaded (long; TimeUnit)` has been added to this interface.

Package Name`org.apache.hadoop.hbase.client`**Effect**

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `setRetryPauseForServerOverloaded (long; TimeUnit)` in `org.apache.hadoop.hbase.client.AsyncAdminBuilder`.

Reason for change[HBASE-26807](#)**Recommendation**

If a class extends `AsyncAdminBuilder` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncBufferedMutator

Abstract method `int getMaxMutations ()` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: a client class `C` is not abstract and does not override abstract method `getMaxMutations ()` in `org.apache.hadoop.hbase.client.AsyncBufferedMutator`.

Reason for change

[HBASE-29148](#)

Recommendation

If a class extends `AsyncBufferedMutator` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncBufferedMutatorBuilder

Abstract method `AsyncBufferedMutatorBuilder setMaxMutations (int)` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: a client class `C` is not abstract and does not override abstract method `setMaxMutations (int)` in `org.apache.hadoop.hbase.client.AsyncBufferedMutatorBuilder`.

Reason for change

[HBASE-29148](#)

Recommendation

If a class extends `AsyncBufferedMutatorBuilder` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncTableBuilder

Abstract method `AsyncTableBuilder<C> setRequestAttribute (String; byte[])` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: a client class `C` is not abstract and does not override abstract method `setRequestAttribute (String; byte[])` in `org.apache.hadoop.hbase.client.AsyncTableBuilder<C>`.

Reason for change

[HBASE-27657](#)

Recommendation

If a class extends `AsyncTableBuilder` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncTableBuilder

Abstract method `AsyncTableBuilder<C> setRetryPauseForServerOverloaded (long; TimeUnit)` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message:
a client class C is not abstract and does not override abstract method
`setRetryPauseForServerOverloaded (long; TimeUnit)` in
`org.apache.hadoop.hbase.client.AsyncTableBuilder<C>`.

Reason for change

[HBASE-26807](#)

Recommendation

If a class extends `AsyncTableBuilder` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

ColumnFamilyDescriptor

Abstract method `IndexBlockEncoding getIndexBlockEncoding ()` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `getIndexBlockEncoding ()` in `org.apache.hadoop.hbase.client.ColumnFamilyDescriptor`.

Reason for change

[HBASE-27314](#)

Recommendation

If a class extends `ColumnFamilyDescriptor` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

ColumnFamilyDescriptor

Abstract method `Compression.Algorithm getMajorCompactionCompressionType ()` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `getMajorCompactionCompressionType ()` in `org.apache.hadoop.hbase.client.ColumnFamilyDescriptor`.

Reason for change

[HBASE-25756](#)**Recommendation**

If a class extends `ColumnFamilyDescriptor` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

ColumnFamilyDescriptor

Abstract method `Compression.Algorithm getMinorCompactionCompressionType ()` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `getMinorCompactionCompressionType ()` in `org.apache.hadoop.hbase.client.ColumnFamilyDescriptor`.

Reason for change

[HBASE-25756](#)

Recommendation

If a class extends `ColumnFamilyDescriptor` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

ServiceCaller

Abstract method `void call (S; RpcController; RpcCallback<R>)` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `call (S; RpcController; RpcCallback<R>)` in `org.apache.hadoop.hbase.client.ServiceCaller<S;R>`.

Reason for change

CDPD-33849

Recommendation

Method signature changed, parameter type(s) changed to HBase thirdparty protobuf class(es). If a client application calls this method, it should import the HBase thirdparty protobuf class(es).

Recompilation Required?

YES

TableBuilder

Abstract method `TableBuilder setRequestAttribute (String; byte[])` has been added to this interface.

Package Name

`org.apache.hadoop.hbase.client`

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `setRequestAttribute (String; byte[])` in `org.apache.hadoop.hbase.client.TableBuilder`.

Reason for change

[HBASE-27657](#)

Recommendation

If a class extends `TableBuilder` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Compression.Algorithm

Abstract method `CompressionCodec reload (Configuration)` has been added to this class.

Package Name

`org.apache.hadoop.hbase.io.compress`

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `reload (Configuration)` in `org.apache.hadoop.hbase.io.compress.Compression.Algorithm`.

Reason for change

[HBASE-26259](#)

Recommendation

If a class extends `Compression.Algorithm` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Removed APIs in 7.3.2

The following APIs are no longer available for HBase in Cloudera Runtime 7.3.2

Table.batchCoprocessorService

Method removed

Package Name

`org.apache.hadoop.hbase.client`

Effect

A client program may be interrupted by `NoSuchMethodError` exception.

Reason for change

CDPD-33849

Recommendation

Method signature changed, return type and parameter types changed to HBase thirdparty protobuf classes. If a client application calls this method, it should import the HBase thirdparty protobuf class(es).

Recompilation Required?

Yes

AuthUtil.getAuthRenewalChore

Method removed

Package Name

`org.apache.hadoop.hbase`

Effect

A client program may be interrupted by `NoSuchMethodError` exception.

Reason for change[HBASE-26212](#)**Recommendation**

Non-public method. Clients should not depend on non-public methods.

Recompilation Required?

No

Bytes.<init>

Method removed

Package Name

org.apache.hadoop.hbase.util

EffectA client program may be interrupted by `NoSuchMethodError` exception.**Reason for change**

None

Recommendation

None

Recompilation Required?

No

PBType.encodedLength

Method removed

Package Name

org.apache.hadoop.hbase.types

EffectA client program may be interrupted by `NoSuchMethodError` exception.**Reason for change**

CDPD-33849

Recommendation

Method signature changed, parameter type changed to HBase thirdparty protobuf class. If a client application calls this method, it should import the HBase thirdparty protobuf class.

Recompilation Required?

Yes

API Compatibility changes in 7.3.2 for Kafka

Removed or Modified APIs in Cloudera Runtime 7.3.2 for Kafka and recommendations for how to handle them.

Apache Version of Kafka in 7.3.1 was 3.4.1 and Apache Version of Kafka in 7.3.2 is 3.9.1

Modified APIs in 7.3.2

The following APIs have been modified for Kafka in Cloudera Runtime 7.3.2

Consumer

Modified

Package Name

org.apache.kafka.clients.consumer

EffectRecompilation of a client program may be terminated with the message, a client class C is not abstract and does not override abstract method `clientId` (Duration) in `org.apache.kafka.clients.consumer`.

Reason for change[KAFKA-15613](#)**Recommendation**

This change affects only the Consumer interface. If you are using the default implementation `KafkaConsumer`, then no change is required. For custom implementation, see [KIP-714: Client metrics and observability](#).

Recompilation Required?

Yes

Producer

Modified

Package Name`org.apache.kafka.clients.producer`**Effect**

Recompilation of a client program may be terminated with the message, a client class `C` is not abstract and does not override abstract method `clientId (Duration)` in `org.apache.Kafka.clients.producer.Producer`.

Reason for change[KAFKA-15613](#)**Recommendation**

This change affects only the Producer interface. If you are using the default implementation `KafkaProducer`, then no change is required. For custom implementation, see [KIP-714: Client metrics and observability](#).

Recompilation Required?

Yes

API Compatibility changes in 7.3.2 for Spark

Removed or Modified APIs in Cloudera Runtime 7.3.2 for Spark and recommendations for how to handle them.

Apache Version of Spark in 7.3.1 was 3.5.4 and Apache Version of Spark in 7.3.2 is 3.5.4

Removed APIs in 7.3.2

The following APIs are no longer available for Spark in Cloudera Runtime 7.3.2

ShuffleExchangeExec.prepareShuffleDependency

Method Removed

Package Name`org.apache.spark.sql.execution.exchange`**Effect**

A client program may be interrupted by `NoSuchMethodError` exception.

Reason for change[SPARK-51016](#)**Recommendation**

None

Recompilation Required?

Yes

API Compatibility changes in 7.3.2 for ZooKeeper

Removed or Modified APIs in Cloudera Runtime 7.3.2 for ZooKeeper and recommendations for how to handle them.

Apache Version of ZooKeeper in 7.3.1 was 3.8.1 and Apache Version of ZooKeeper in 7.3.2 is 3.8.5

Removed APIs in 7.3.2

The following APIs are no longer available for ZooKeeper in Cloudera Runtime 7.3.2

SecurityUtils.createSaslClient

Method Removed

Package Name

org.apache.zookeeper.util

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[ZOOKEEPER-4889](#)

Recommendation

Client code need to be adjusted to supply ZooKeeper Configuration to the method.

Recompilation Required?

Yes

Deprecation Notices In Cloudera Runtime 7.3.2.0

Components and features that will be deprecated or removed in this release or a future release.

Terminology

Items in this section are designated as follows:

Deprecated

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

Moving

Technology that Cloudera is moving from a future Cloudera 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 release and plan for the alternative Cloudera offering or subscription for the technology.

Removed

Technology that Cloudera has removed from Cloudera 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.

Removed Components and Product Capabilities

- Apache Spark 2

Spark 3 is the default Spark version in Cloudera Runtime. Spark 2 (and Livy 2) has been removed and no longer available in 7.3.1

**Important:**

Spark 3 contains a large number of changes from Spark 2.

Refer to *Upgrading Spark* for more information on upgrading Spark clusters to 7.3.1.0, and *Migrating Spark Applications* for more information on migrating your existing Spark applications between versions 2 and 3.

- Apache Livy 2 (see [Deprecation Notices for Apache Livy](#))
- Apache Zeppelin (see [Deprecation Notices for Apache Zeppelin](#))

Please contact Cloudera Support or your Cloudera Account Team if you have any questions.

Platform and OS

Cloudera Runtime 7.3.2 removes support for several operating systems, databases, Python versions, and Java Development Kits (JDKs). To ensure a successful upgrade, verify that your environment meets these new requirements.

Database Support:

Starting with Cloudera Runtime 7.3.2, Cloudera no longer supports the following databases:

- PostgreSQL 13
- Oracle 19c (including 19c RAC)

Operating System

Starting with Cloudera Runtime 7.3.2, Cloudera no longer supports the following operating systems:

- SLES 15 SP4
- Ubuntu 20.04

Python Support

Cloudera Runtime 7.3.2 limits Python compatibility to modern versions, and now only supports Python 3.11. Consequently, Cloudera has:

- Removed support for Python 3.8
- Removed support for Python 3.9
- Removed support for Python 3.10

JDK Support

To improve performance and security, Cloudera Runtime 7.3.2 consolidates Java support, and now only supports JDK 17. Consequently, Cloudera has:

- Removed support for Azul JDK 8, Open JDK 8, and Oracle JDK 8
- Removed support for Azul JDK 11, Open JDK 11, and Oracle JDK 11

Deprecation Notices for Cruise Control

Certain features and functionality in Cruise Control are deprecated or removed in Cloudera Runtime 7.3.2.

Removed

Cloudera Manager metrics reporter

Cloudera Manager metrics reporter is removed from Cloudera Runtime 7.3.2. Migrate to Cruise Control metrics reporter.

Deprecation Notices for Apache Ranger

Certain features and functionality in Apache Ranger are deprecated or removed in Cloudera Runtime 7.3.2.0.

Removed

Deprecation of Ranger Backbone Classic UI

Starting from Cloudera Runtime 7.3.2.0, the Backbone Classic UI option has been removed from the Ranger Admin interface. Support for the legacy Backbone-based UI is no longer available.

Deprecation Notices for Apache Kafka

Certain features and functionality in Apache Kafka are deprecated or removed in Cloudera Runtime 7.3.2. You must review these changes along with the information about the features in Kafka that will be removed or deprecated in a future release.



Important: The following list of deprecated and removed items is not exhaustive and only contains items that have a direct and immediate effect on Kafka in Cloudera. For a full list of deprecation and/or removals in the version Apache Kafka shipped with Cloudera Runtime, review the *Notable Changes* as well as the *Release Notes* on <https://kafka.apache.org/>.

Deprecated

ZooKeeper

Deploying new or using existing Kafka clusters running in ZooKeeper mode is deprecated. Support for ZooKeeper-based Kafka clusters will be removed in a future release. Cloudera recommends the following:

- Deploy all new Kafka clusters in KRaft mode.
- Migrate existing ZooKeeper-based clusters to KRaft following an upgrade to Cloudera Runtime 7.3.2.

This is the only version where migration is possible. Neither previous or future major, minor, and maintenance versions support migration.

For additional information, see the following resources:

- [Installing Kafka in Cloudera Base on premises](#)
- [Kafka KRaft overview](#)
- [Migrating Kafka from ZooKeeper to KRaft overview](#)

MirrorMaker (MM1)

MirrorMaker is deprecated. Cloudera recommends that you use Streams Replication Manager instead.

--zookeeper

The --zookeeper option is only supported for the kafka-configs tool and should be only used when updating SCRAM Credential configurations. The --zookeeper option is either deprecated in or removed from other Kafka command line tools. Cloudera recommends that you use the --bootstrap-server option instead.

Deprecation Notices for Streams Replication Manager

Certain features and functionality in Streams Replication Manager are deprecated or removed in Cloudera Runtime 7.3.2.

Deprecated

Legacy per-flow REST server implementation

The legacy per-flow REST server implementation of Streams Replication Manager is deprecated in Cloudera Runtime 7.3.2 and will be removed in a future release. In the previous implementation, the Streams Replication Manager Driver role started a dedicated REST server for each replication flow. The new implementation uses a single REST server for all flows. Cloudera recommends that you

migrate your Streams Replication Manager clusters to the new implementation as part of upgrading to 7.3.2. For migration steps, see [Configuring the SRM REST server for replication](#).