

Cloudera Runtime 7.3.1

Release Notes

Date published: 2020-07-28

Date modified: 2024-12-10

CLOUDERA

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

Legal Notice

© Cloudera Inc. 2024. All rights reserved.

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

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

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

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

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

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

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

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

Contents

Overview.....	6
What's New In Cloudera Runtime 7.3.1.....	6
Atlas.....	7
What's New in Cloud Connectors.....	9
Cruise Control.....	9
Hive.....	9
Hue.....	12
Iceberg.....	12
Impala.....	13
Kafka.....	15
Kerberos.....	16
Livy.....	16
Navigator Encrypt.....	16
Oozie.....	16
Ranger.....	17
Ranger KMS.....	18
Schema Registry.....	18
Solr.....	18
Spark.....	18
SMM.....	19
SRM.....	20
YARN and YARN Queue Manager.....	21
Unaffected Components in this release.....	22
What's new in Platform Support.....	22
Cloudera Runtime Component Versions.....	22
Using the Cloudera Runtime Maven repository 7.3.1.....	24
Runtime 7.3.1.0-197.....	24
Fixed Issues In Cloudera Runtime 7.3.1.....	42
Atlas.....	42
Avro.....	43
Cloud Connectors.....	43
Cruise Control.....	43
Hadoop.....	43
HDFS.....	44
HBase.....	44
Hive.....	47
Hue.....	51
Impala.....	52
Fixed Issues in Apache Iceberg.....	52

Kafka.....	52
Kudu.....	53
Knox.....	53
Livy.....	55
Navigator Encrypt.....	55
Oozie.....	56
Ozone.....	56
Parquet.....	59
Phoenix.....	60
Ranger.....	60
Ranger KMS.....	64
Schema Registry.....	65
Solr.....	65
Spark.....	66
SMM.....	66
SRM.....	66
Tez.....	67
YARN and YARN Queue Manager.....	67
Zookeeper.....	67

Known Issues In Cloudera Runtime 7.3.1..... 67

Atlas.....	67
Avro.....	69
Known Issues in Cloud Connectors.....	69
Cruise Control.....	70
Hadoop.....	70
HBase.....	70
HDFS.....	72
Hive.....	75
Hue.....	75
Iceberg.....	79
Impala.....	80
Kafka.....	84
Kerberos.....	88
Knox.....	88
Kudu.....	89
Navigator Encrypt.....	90
Oozie.....	90
Ozone.....	91
Parquet.....	99
Phoenix.....	99
Ranger.....	99
Ranger KMS.....	101
Schema Registry.....	101
Solr.....	102
Spark.....	110
Sqoop.....	110
SMM.....	111
SRM.....	112
MapReduce, YARN and YARN Queue Manager.....	113
Zookeeper.....	119

Behavioral Changes In Cloudera Runtime 7.3.1..... 120

Behavioral Changes in Atlas.....	120
----------------------------------	-----

Behavioral Changes in Hive.....	120
Behavioral Changes in Impala.....	121
Behavioral Changes in Knox.....	121
Behavioral Changes in Livy.....	122
Behavioral Changes in Ranger.....	122
Behavioral Changes in Spark.....	123
Deprecation Notices In Cloudera Runtime 7.3.1.....	124
Platform and OS.....	124
Deprecation Notices for Apache Kafka.....	125
Deprecation Notices for Apache Livy.....	125
Navigator Key Trustee Server.....	126
Deprecation Notices for Apache Oozie.....	126
Deprecation Notices for Apache Spark.....	126
Deprecation Notices for Apache Zeppelin.....	127
Fixed Common Vulnerabilities and Exposures 7.3.1.....	127

Overview

This document provides you with the latest information about Cloudera Runtime 7.3.1. It includes improvements and describes new features, bug fixes, tech previews and more. For detailed information about the runtime components, see [Cloudera documentation](#).

What's New In Cloudera Runtime 7.3.1

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

Spark 2 removed from Cloudera Runtime

Spark 3 is the default Spark version in Cloudera Runtime. Spark 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, and *Migrating Spark Applications* for more information on migrating your existing Spark applications between versions 2 and 3.

Upgrade CDP 7.1.7 SP3 to CDP 7.3.1

You can perform an In-place upgrade from CDP 7.1.7 SP3 to CDP 7.3.1. For more information, see [CDP to CDP](#) documentation.

Upgrade CDP 7.1.8 CHF and below to CDP 7.3.1

You can perform an In-place upgrade from CDP 7.1.8 CHF and below to CDP 7.3.1. For more information, see [CDP to CDP](#) documentation.

Upgrade CDP 7.1.9 to CDP 7.3.1

You can perform an In-place upgrade from CDP 7.1.9 to CDP 7.3.1. For more information, see [CDP to CDP](#) documentation.

Upgrade CDP 7.1.9 SP1 to CDP 7.3.1

You can perform an In-place upgrade from CDP 7.1.9 SP1 to CDP 7.3.1. For more information, see [CDP to CDP](#) documentation.

Rollback CDP 7.3.1. to CDP 7.1.7 SP3

You can downgrade or rollback an upgrade from CDP Private Cloud Base 7.3.1 to CDP 7.1.7 SP3. The rollback restores your CDP cluster to the state it was in before the upgrade, including Kerberos and TLS/SSL configurations. For more information, see [Rollback CDP 7.3.1 to CDP 7.1.7 SP3](#) documentation.

Rollback CDP 7.3.1 to CDP 7.1.8 CHF

You can downgrade or rollback an upgrade from CDP Private Cloud Base 7.3.1 to CDP 7.1.8 CHF. The rollback restores your CDP cluster to the state it was in before the upgrade, including the Kerberos and TLS/SSL configurations. For more information, see [Rollback CDP 7.3.1 to CDP 7.1.8 CHF](#) documentation.

Rollback CDP 7.3.1 to CDP 7.1.9

You can downgrade or rollback an upgrade from CDP Private Cloud Base 7.3.1 to CDP 7.1.9. The rollback restores your CDH cluster to the state it was in before the upgrade, including the Kerberos and TLS/SSL configurations. For more information, see [Rollback CDP 7.3.1 to CDP 7.1.9](#) documentation.

Rollback CDP 7.3.1 to CDP 7.1.9 SP1

You can downgrade or rollback an upgrade from CDP Private Cloud Base 7.3.1 to CDP 7.1.9 SP1. The rollback restores your CDH cluster to the state it was in before the upgrade, including the Kerberos and TLS/SSL configurations. For more information, see [Rollback CDP 7.3.1 to CDP 7.1.9 SP1](#) documentation.

Related Information

[Upgrading Spark](#)

[Migrating Spark Applications](#)

What's New in Apache Atlas

Learn about the new features of Apache Atlas in Cloudera Runtime 7.3.1.

Cloudera Navigator Data Management upgrade to Apache Atlas

In CDP, Apache Atlas fulfills the metadata collection role that in CDH was filled by Cloudera Navigator Data Management. The upgrade to CDP provides a method to migrate Navigator content, including technical and business metadata, to Atlas. For information, see [Migrating Navigator content to Atlas](#). For the cluster audit functionality handled by Navigator, see the (production version of) access auditing provided by Apache Ranger see [Ranger Audit Overview](#).

Business Metadata: Entity model extensions

This release of Atlas provides the ability for data stewards to add custom attributes to existing entity types and set their values on existing entities. This functionality allows an organization to extend its enterprise data model with curated master data attributes that have specific meaning for the business. Business Metadata attributes are defined centrally and can be used on designated entity types. Administrators can control who can view, add values to, and create or update set collections of Business Metadata attributes. Privileged users can add free-form values or select from predefined values to the attribute for a given entity. For more information, see [Leveraging Business Metadata](#).

Bulk import of Business Metadata attribute associations

Atlas provides an interface to import a list of assignments of Business Metadata attributes to entities. The list includes information to uniquely identify the Business Metadata attribute and the targeted entity. The list can be formatted as comma-separated values (.CSV) or Microsoft Excel (.XLS) formatted file. For more information, see [Importing Business Metadata associations in bulk](#)

Bulk import of Glossary terms

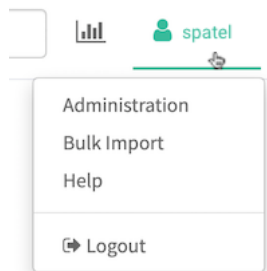
Atlas provides an interface to import a list of terms into existing Glossaries. The list can include any or all of the metadata associated with a given term. The list can be formatted as comma-separated values (.CSV) or Microsoft Excel (.XLS) formatted file. For more information, see [Importing Glossary terms in bulk](#)

Administrator features have a home in the Atlas UI

The Atlas UI now contains an Administration section available to users with administrator privileges:

- Review system-level audits, such as created by entity purge events. See [Auditing purged entities](#).
- Create enumerations for use as attribute values. See [Defining Apache Atlas enumerations](#).
- Create Business Metadata attributes. See [Creating Business Metadata](#).

Open the Administration section from the user menu at the top right of the Atlas UI.



Purge of deleted entities

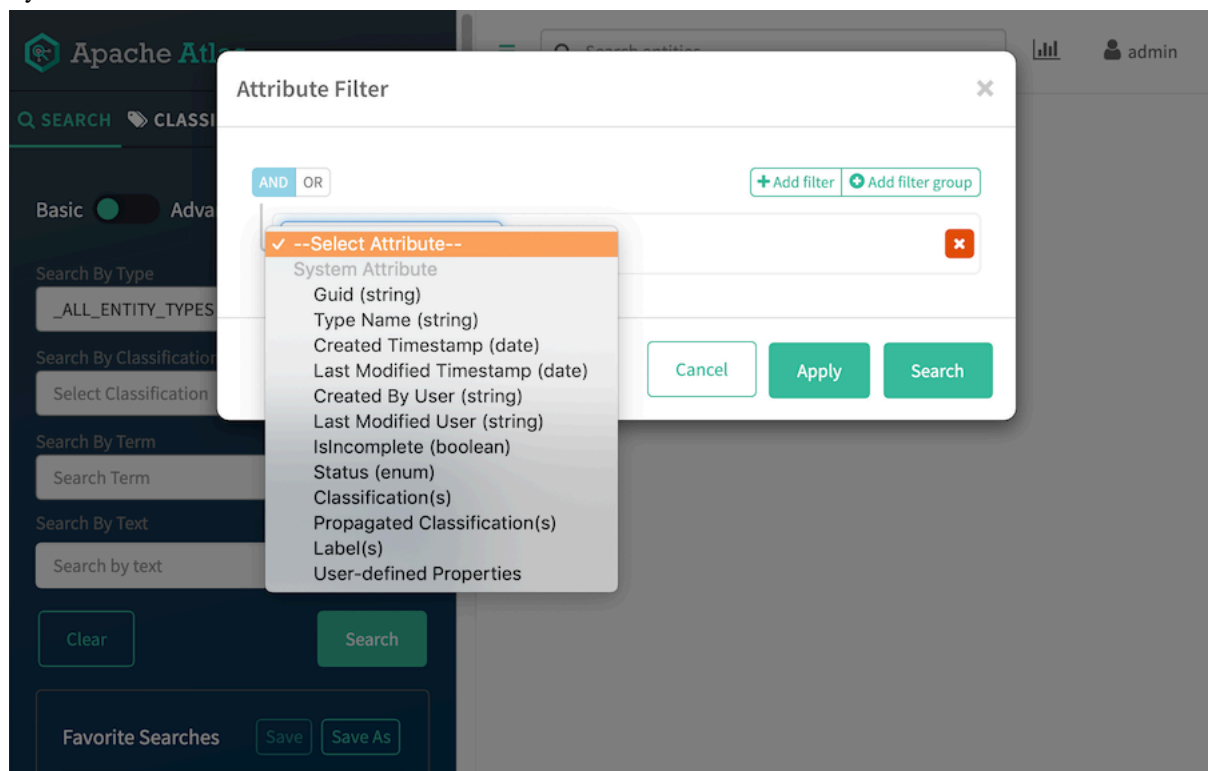
Atlas now provides the ability to clear the metadata for entities that represent data assets and operation that no longer exist on the cluster. The purge functionality is available to users with administrator privilege; run a REST API command that lists one or more GUID values for the deleted entities. For more information, see [Purging deleted entities](#).

Enhancements to Basic Search in Atlas

Atlas Basic Search includes a filter to allow users to search for entities based on values of entity attributes. In this release, the search filter includes access to system attributes, labels, classifications, and user-defined properties. The filter allows users to build logical combinations of search criteria, including multiple classifications. For more information, see [Using Basic Search](#).

System attributes filter searches

Atlas basic and advanced search now allow you to filter based on system attribute values, including when and by whom an entity was created. Classifications are also modeled as system attributes, so this change allows you to filter on the classifications assigned to an entity and to distinguish between classifications and propagated classifications. System attributes are available in the search filter.



For information on using system attributes in Advanced Search, see [Apache Atlas metadata attributes](#).

What's New in Cloud Connectors

Learn about the new features of Cloud Connectors in Cloudera Runtime 7.3.1.

Support for Amazon S3 Express One Zone Storage

Support for Amazon S3 Express One Zone is added. The Amazon S3 Express One Zone is single-Availability Zone and high-performance storage class that delivers consistent single-digit millisecond data access. A specific Availability Zone can be selected within an AWS Region to store your data. This enables you to have your storage and compute resources in the same Availability Zone to optimize performance, lower compute costs and run workloads faster. Your data is stored in an S3 directory bucket that supports hundreds of thousands of requests per second.

Vectorized IO support

Support for Hadoop Vectorized IO API is added for ORC and Parquet file formats. The S3A connector offers a customized implementation that enables parallel and asynchronous reading of different data blocks.

Migration to AWS V2 SDK

AWS V2 Java SDK is used for communicating with AWS services, which includes storage through the S3A connector.

The following improvements and enhancements have been introduced for AWS V2 SDK migration:


- Dual-layer server-side encryption (DSSE) has been enabled with AWS KMS keys
- AWS SDK V2 has been upgraded to 2.25.53

What's New in Cruise Control

Learn about the new features of Cruise Control in Cloudera Runtime 7.3.1.

Cruise Control is added to Streams Messaging Manager UI

A new page is added to Streams Messaging Manager to monitor the Kafka cluster state and rebalancing process with Cruise Control. The Cruise Control User Interface (UI) enables you to review and configure the rebalancing of Kafka clusters through dashboards and a rebalancing wizard. The available goals and anomaly detectors are based on the

Cloudera Manager configurations of Cruise Control. You can access Cruise Control from SMM using the  on the navigation sidebar.

For more information about Cruise Control in SMM, see [Monitoring and managing Kafka cluster rebalancing](#).

What's New in Apache Hive

Learn about the new features of Hive in Cloudera Runtime 7.3.1.

Trusting HTTP headers for authentication

When HTTP headers are authenticated via Knox, they can be trusted to establish a session without re-authenticating at HiveServer2. If a trusted header is present in the HTTP request, password-based authentication is skipped, and the client name is extracted directly from the Authorization header.

This change simplifies the authentication process, eliminating the need for repeated authorization since the trusted header confirms the user has already been authenticated by Knox.

Apache Jira : [HIVE-25349](#)

Multi-authentication support for SAML and LDAP in Hive

You can now connect to Hive using both SAML and LDAP authentication modes simultaneously when the transport mode is set to HTTP. This new feature allows you to use multiple authentication mechanisms concurrently, making it easier to connect without adjusting the authentication settings for different use cases. The configuration `hive.server2.authentication` now accepts comma-separated values for both SAML and LDAP.

Apache Jira : [HIVE-25875](#)

Improved query plans using constraint information

Hive now uses constraint information, such as not null, when creating RexNodes, leading to more optimized query plans. This update enables Hive to generate simpler, more efficient query plans by avoiding unnecessary joins when not null constraints are applied.

Apache Jira : [HIVE-26043](#)

Print DAG ID to console

You can now view the DAG ID directly in the console when executing queries. This makes it easier to track and debug query executions by providing immediate visibility of the DAG ID.

Apache Jira : [HIVE-25176](#)

Increase default value of PartitionManagementTask frequency

The default frequency for `metastore.partition.management.task.frequency` has been increased from five minutes to six hours. This change ensures better performance in production environments with a lot of databases and tables, allowing enough time for the task to scan all tables and partitions.

Apache Jira : [HIVE-27011](#)

Support for both LDAP and kerberos authentication in HiveServer2

HiveServer2 now supports both LDAP and kerberos authentication simultaneously. The configuration `hive.server2.authentication` can accept comma-separated values for both Kerberos and LDAP even in binary mode.

Apache Jira: [HIVE-27352](#)

Thrift-over-HTTP support for Hive Metastore client

Hive Metastore client can now connect through Thrift-over-HTTP, enabling access through Knox.

Apache Jira: [HIVE-21456](#)

Data connector authorization on the Hive Metastore server side

You can now authorize Data Definition Language (DDL) operations for connectors on the Hive Metastore server side. This enhancement improves security by ensuring only authorized users can perform these operations.

Apache Jira: [HIVE-26248](#)

Setting the user for compaction tasks

This update introduces a new configuration that allows you to specify a user for running compaction tasks, instead of relying on the table directory owner by default. The configuration provides flexibility by enabling you to assign a specific user for compaction operations, including file listing in the Initiator and Cleaner.

This is useful when you need to run compaction as a specific user, giving administrators more control over permissions and task management. Admins can now optionally configure the user that will handle compaction tasks.

Apache Jira: [HIVE-24191](#)

Support for HDFS snapshots

You can now utilize HDFS snapshots to enhance external table replication. With the addition of DistCp diff using snapshots, replication is streamlined to include only modified entries. This eliminates the need to list all files and directories, significantly reducing the effort and time required for data copying.

Apache Jira: [HIVE-24852](#)

Ability to create tables on individual files directly

You can now create tables directly on individual files within a directory in Hive. This feature allows you to define tables for specific files without changing the existing directory structure, enabling seamless data management for multiple teams using a common directory.

Apache Jira: [HIVE-25569](#)

New API for retrieving all table constraints

You can now use the `getAllTableConstraints` API to retrieve all table constraints such as Primary Key, Foreign Key, and others in a single call. This improvement consolidates multiple metastore calls into one, reducing the need for separate requests and improving efficiency. Local caching is also added to HiveServer to avoid duplicate calls to Hive Metastore.

Apache Jira: [HIVE-22782](#)

Beeline standalone execution with Java

You can now run Beeline as a standalone tool using Java without relying on `HADOOP_HOME`. A new distributable tarball isolates all necessary dependencies, allowing Beeline to run with just JRE and the required jars. This simplifies execution on edge nodes without needing a full Hive or Hadoop setup.

Apache Jira: [HIVE-24348](#)

JWT authentication support in HTTP mode

You can now use JWT for authentication in HiveServer when running in HTTP mode. HiveServer retrieves the JWKS and verifies the JWT in the Authorization header, while the JDBC client can accept JWTs from either the environment variable or the JDBC URL, sending it in the Authorization header.

Apache Jira: [HIVE-25575](#)

Vectorization support for lead and lag functions

You can now benefit from vectorized execution for lead and lag functions, improving performance through better vectorization coverage.

Apache Jira: [HIVE-24945](#)

Dynamic connection pool for TxnHandler#connPoolMutex

You can now benefit from a dynamic connection pool for `TxnHandler#connPoolMutex`, replacing the fixed-size pool. This change allows the pool to scale by adding or closing connections on demand, improving resource efficiency for non-leader instances in the warehouse and making the Hive Metastore more scalable.

Apache Jira: [HIVE-26794](#)

Upgrade ORC to version 1.8.3

Hive now supports ORC version 1.8.3, offering improved memory usage and performance.

Apache Jira: [HIVE-26809](#)

Support for generic LDAP search bind filters in Hive

Hive's LDAP authentication has been enhanced to support generic LDAP search bind filters, making it easier to configure. New configurations have been added:

- `hive.server2.authentication.ldap.userSearchFilter`
- `hive.server2.authentication.ldap.groupSearchFilter`
- `hive.server2.authentication.ldap.groupBaseDN`

These configurations will work alongside the existing `hive.server2.authentication.ldap.baseDN`. You can choose to use these new options or continue with the current setup, ensuring backward compatibility.

Apache Jira: [HIVE-27311](#)

What's New in Hue

Learn about the new features of Hue in Cloudera Runtime 7.3.1.

General availability (GA) of the SQL AI Assistant

Hue leverages the power of Large Language Models (LLM) to help you generate SQL queries from natural language prompts and also provides options to optimize, explain, and fix queries, promoting efficient and accurate practices for accessing and manipulating data. You can use several AI services and models such as OpenAI's GPT service, Amazon Bedrock, and Azure's OpenAI service to run the Hue SQL AI assistant.

- To learn more about the supported models and services, limitations, and what data is shared with the LLMs, see [About the Hue SQL AI Assistant](#).
- To set up and enable the SQL AI Assistant, see [About setting up the Hue SQL AI Assistant](#).
- To see how to generate, edit, explain, optimize, and fix queries, see [Starting the SQL AI Assistant in Hue](#).

Hue supports Python 3.9 on RHEL 8.8 and RHEL 8.10

Starting from the 7.3.1 release, Hue supports only Python 3.9 for RHEL 8.8 and RHEL 8.10. Before upgrading to CDP runtime 7.3.1, you must install Python 3.9 on all the Hue servers, as Hue requires a Python 3.9 version and does not start without it. For information about migrating from Python 3.8 to Python 3.9, see [Migrating from Python 3.8 to Python 3.9 on RHEL 8.8 or RHEL 8.10](#).

What's New in Apache Iceberg

Learn about the new features of Iceberg in Cloudera Runtime 7.3.1.

Apache Iceberg support for Hive

Cloudera Data Platform (CDP) supports a Data Lakehouse architecture by pre-integrating and unifying the capabilities of Data Warehouses and Data Lakes, to support data engineering, business intelligence, and machine learning – all on a single platform.

Starting from this release, CDP Private Cloud Base supports queries of Iceberg tables from the Apache Hive compute engine. You can run SQL queries to create and query Iceberg tables. Hive queries are table-format agnostic. You can run nested, correlated, or analytic queries on all supported table types. Hive on Iceberg supports and enables you to use the following Apache Iceberg features:

- ACID transactions with Iceberg V2 tables

- Point in time queries using Iceberg Time travel
- Rollback table
- Position deletes
- Schema evolution
- Flexible partitioning using partition evolution and partition transform
- Support for materialized views
- Snapshot expiry
- Merge table
- Multi-engine concurrent read and write

For more information about the Apache Iceberg features supported in CDP, see [Using Apache Iceberg](#).

If you want to migrate your existing Hive tables to Iceberg tables, you can use the ALTER TABLE statement. For more information, see [Migrate Hive table to Iceberg](#).

CDP supports the [integration of Iceberg and Atlas](#) that helps you identify the Iceberg tables to scan data and provide lineage support. Learn how Atlas works with Iceberg and what schema evolution, partition specification, partition evolution are with examples.

What's New in Apache Impala

Learn about the new features of Apache Impala in Cloudera Runtime 7.3.1.

Added read support for PageHeaderV2 to the Parquet scanner

This update moves page reading logic to new classes, ParquetColumnChunkReader and ParquetPageReader, simplifying V2 data page reading and decompression. It enhances code manageability for both V1 and V2 formats.

Apache Jira: [IMPALA-6433](#)

Collections of fixed length types as non-passthrough children of unions

This update enables collections of fixed-length types to be used as non-passthrough children in UNION ALL operations. It achieves this by allowing the materialization of these collections.

Apache Jira: [IMPALA-12147](#)

Display query execution progress in Impala Web UI

Adds a query progress indicator to the /queries page in Impala's Web UI, showing the completion status of fragment instances. This feature provides better tracking for computation-intensive queries, supplementing the scan progress bar.

Apache Jira: [IMPALA-12048](#)

Allow implicit casts between numeric and string types when inserting into table

The current implementation requires explicit casts for numeric and string-based literals. This is controlled through a query option `allow_unsafe_casts` and turned off by default. This query option allows implicit casting between some numeric types and string types. See, [implicit casting](#)

Apache Jira: <https://issues.apache.org/jira/browse/IMPALA-10173>

Optimize query planning by reducing getLocation() and getFileSystem() calls

The fix reduces planning time by calling HdfsPartition.getLocation() once per partition and caching the FileSystem object based on the URI scheme and authority. This minimizes expensive decompression and redundant getFileSystem() calls, improving performance for queries with many partitions.

Apache Jira: [IMPALA-12408](#)

JSON File Reader Prototype

This prototype enables reading JSON files using the rapidjson library with Arrow support such as HdfsJsonScanner, callback functions, and startup flag.

Apache Jira: [IMPALA-10798](#)

CREATE TABLE LIKE for Kudu tables

Impala now supports a dedicated keytab for HTTP SPNEGO authentication, enabling easier management of Kerberos keytabs. A new `--spnego_keytab_file` flag lets you specify a separate keytab for the web console when `--webserver_require_spnego` is enabled. If this flag is set, the web server will use the SPNEGO keytab for HTTP authentication, while the main service keytab remains unchanged. If not specified, the web server defaults to using the primary service keytab for SPNEGO

Apache Jira: [IMPALA-4052](#)

Dedicated SPNEGO keytab for Impala web console authentication

Impala now supports a dedicated keytab for HTTP SPNEGO authentication, enabling easier management of Kerberos keytabs. A new `--spnego_keytab_file` flag lets you specify a separate keytab for the web console when `--webserver_require_spnego` is enabled. If this flag is set, the web server will use the SPNEGO keytab for HTTP authentication, while the main service keytab remains unchanged. If not specified, the web server defaults to using the primary service keytab for SPNEGO

Apache Jira: [IMPALA-12318](#)

Non-unique primary keys in Kudu

Kudu now supports non-unique primary keys by automatically adding an `auto_increment_id` column to form a unique composite primary key. This column, a system-generated big integer, ensures uniqueness within each tablet server region and is hidden unless specified in `SELECT` statements. `ALTER TABLE` modifications and `UPSERT` operations for this column are currently unsupported.

Apache Jira: [IMPALA-11809](#)

Hive's ESRI geospatial functions as built-ins

This change adds Hive's ESRI geospatial functions as built-in UDFs in Impala.

Apache Jira: [IMPALA-11745](#)

Unicode column name support in Impala

Impala now supports Unicode characters in column names, aligning with Hive's support for non-ASCII characters. This enhancement leverages Hive's `validateColumnName()` function, which removes restrictions on column names at the metadata level. With this update, Impala allows greater flexibility for column naming while remaining consistent with Hive's metadata validation standards.

Apache Jira: [IMPALA-12465](#)

Support custom hash partitions at range level in Kudu tables

Impala now supports specifying custom hash partitions at the range level in Kudu tables. You can define hash schemas within specific partitions using the updated `CREATE TABLE` and `ALTER TABLE` syntax, and view them with the new `SHOW HASH SCHEMA` statement. This update aligns hash partitioning more closely with range partitioning, enhancing flexibility while maintaining backward compatibility.

Apache Jira: [IMPALA-11430](#)

What's New in Apache Kafka

Learn about the new features of Apache Kafka in Cloudera Runtime 7.3.1.

Kafka Rolling Restart check—all partitions fully replicated

A new broker rolling restart check option, all partitions fully replicated has been introduced. Selecting this option ensures that all partitions are in a fully synchronized state when a broker is stopped. For more information, see [Rolling restart checks](#).

Kafka is safely stopped during operating system upgrades

During OS upgrades, Cloudera Manager now ensures that Kafka brokers are safely stopped. Specifically, Cloudera Manager now performs a rolling restart check before stopping a broker. This ensures that the Kafka service stays healthy during the upgrade. The level of health guarantee that Cloudera Manager ensures is determined by the restart check type set in the Cluster Health Guarantee During Rolling Restart Kafka property. Cloudera recommends that you set this property to all partitions stay healthy to avoid service outages. For more information, see [Rolling restart checks](#).

useSubjectCredsOnly set to true by default in Kafka Connect

In previous versions, the `javax.security.auth.useSubjectCredsOnly` JVM property was set to false in Kafka Connect. Because of this, connectors running with an invalid or no JAAS configuration could use the credentials of other connectors to establish connections. Starting with this release, `useSubjectCredsOnly` is set to true by default. As a result, connectors are required to use their own credentials.

This default change is true for newly provisioned clusters. On upgraded clusters, `useSubjectCredsOnly` remains set to false to ensure backwards compatibility. If you are migrating connectors from a cluster running a previous version of Runtime to a new cluster running 7.2.18 or later, you must ensure that credentials are added to the connector configuration when migrated. Otherwise, migrated connectors may not work on the new cluster.

In addition to the default value change, a new Kafka Connect property is introduced in Cloudera Manager that you can use to set `useSubjectCredsOnly`. The property is called Add Use Subject Credentials Only JVM Option With True Value. Setting this property to false does not expressly set `useSubjectCredsOnly` to false. Instead, it sets `useSubjectCredsOnly` to the cluster default value.

Persistent MQTT sessions support for the MQTT Source connector

Version 1.1.0 of the MQTT Source connector is released. The connector now supports MQTT persistent sessions. This enables the connector to resume (persist) a previous session with an MQTT broker after a session is interrupted. Enabling this feature can ensure that no messages are lost if the connector is momentarily stopped or if the network connection is interrupted.

To support persistent sessions, the following new properties are introduced:

- MQTT Client ID

This property specifies the MQTT client ID that the connector uses.

- MQTT Clean Session

This property controls whether the connector should start clean or persistent sessions. Set this property to false to enable persistent sessions.

Existing connectors will continue to function, upgrading them, however, is not possible. If you want to use the new version of the connector, you must deploy a new instance of the connector. For more information, see [MQTT Source connector](#) and [MQTT Source properties reference](#).

What's New in Kerberos

Learn about the new features of Kerberos in Cloudera Runtime 7.3.1.

Support for Custom Kerberos Principals and System Users

Cloudera Manager configures CDP services to use the default Kerberos principal names and default System Users. While it is possible to customize the Kerberos principal names or System Users for most cluster services by setting various configuration properties, it requires extensive custom configuration. If your security policies require you to customize the service Kerberos Principals and System User Names, Cloudera recommends working closely with Cloudera Professional services in doing so. For more information, see [Customizing Kerberos Principals and System Users](#).

What's New in Livy

Learn about the new features of Livy in Cloudera Runtime 7.3.1.

High Availability support added for Livy

Livy now supports high availability. If there are more than one Livy Servers in the cluster, high availability is automatically enabled.

What's New in Navigator Encrypt

Learn about the new features of Navigator Encrypt in CDP Private Cloud Base 7.3.1.

Migration of Navigator Encrypt keys with HSM from Key Trustee Server (KTS) to Ranger KMS

As Ranger KMS KTS and Key Trustee Server are not supported anymore, you need to migrate your Navigator Encrypt keys with HSM from Ranger KMS KTS and Key Trustee Server to Ranger KMS. For more information, see [Navigator Encrypt key migration with HSM](#).

Navigator Encrypt now supports LUKS2

Navigator Encrypt now supports LUKS2. LUKS implements a platform-independent standard on-disk format for use in various tools. Version 2 of LUKS supports various new features. For details, see the LUKS documentation. NavEncrypt defaults to use whichever version of LUKS the platform uses, for example, on RHEL 9 its LUKS2. For more information, see [LUKS2 support](#).

What's New in Apache Oozie

Learn about the new features of Apache Oozie in Cloudera Runtime 7.3.1.

OpenJPA 3 upgrade in Apache Oozie

Third-party library OpenJPA is upgraded from version 2 to version 3 in Apache Oozie. This upgrade includes the following updates:

- New configuration properties
- Deprecated configuration properties
- Enhanced error handling

For more information, see [OpenJPA upgrade](#).

Updating column types in Oozie Database

If the Oracle database is used for Oozie, then you must update the APP_PATH column type to store values with more than 255 characters. This ensures Oozie does not get stuck in PREP state when your application path exceeds the 255 character limit. If the APP_PATH column type is not

updated, then Oozie fails to run the jobs with the following database error message Data too long for column 'app_path'. This scenario is also applicable to coordinator and bundle jobs. For database types other than Oracle, this update is not mandatory for using Oozie. It works without the update if the APP_PATH value does not exceed 255 characters. Also, Oozie's internal database schema validation fails with an unexpected APP_PATH column type. However, this validation will not have any effect. It just logs its result.

For more information about how to update column types in Oozie Database, see [Updating column types in Oozie Database](#).

What's New in Apache Ranger

Learn about the new features of Apache Ranger in CDP Private Cloud Base 7.3.1.

Support multiple columns policy creation in Ranger for Grant/Revoke request

This enhancement supports multiple columns policy creation in Ranger for Grant/Revoke requests for Impala.

Ranger REST API improvements

Ranger REST APIs have the following changes:

- The following APIs have been removed:
 - assets/credstores - GET, POST, PUT
 - credstores/count - GET
 - credstores/{id} - GET
 - /xusers/auditmaps - GET
 - /xusers/auditmaps/count - GET
 - /xusers/permmaps - GET
 - /resource/{id} - GET
 - assets/policyList/{repository}
 - /groupgroups/* (All methods)
- The following APIs were not returning any access code when request is denied; now they suppose to 403:
 - /tags/tags
 - /tags/types
 - /tags/resources APIs
- Earlier When a non admin user makes a DELETE request to below endpoint, it was returning 405 method not allowed. However, now it returns 403.
 - /assets/resources/{resource_id}
- Earlier the API was not accessible for the keyadmin role users, but now it shall be accessible.
 - /xaudit/trx_log
- Earlier the below mentioned API was returning {OWNER} and {USER} users in the response but now onwards it will not return because access to the users list will be based on which role user is having permissions to which role user.
 - /service/xusers/users
- The API endpoint /xaudit/trx_log/{trx_log_id} was not accessible by keyadmin users. keyadmin users can access the transaction logs using the endpoint /xaudit/trx_log, hence, the keyadmin users should also be allowed to access the endpoint /xaudit/trx_log/{trx_log_id} for transaction log ids related to KMS audits.

Ranger KMS

Learn about the new features of Ranger KMS in CDP Private Cloud Base 7.3.1.

Migration of keys from Key Trustee Server (KTS) to Ranger KMS

As Ranger KMS KTS and Key Trustee Server are not supported anymore, you need to migrate your encryption keys from Ranger KMS KTS and Key Trustee Server to Ranger KMS. For more information, see [Key migration in UCL](#).

Migration of Navigator Encrypt keys with HSM from Key Trustee Server (KTS) to Ranger KMS

As Ranger KMS KTS and Key Trustee Server are not supported anymore, you need to migrate your Navigator Encrypt keys with HSM from Ranger KMS KTS and Key Trustee Server to Ranger KMS. For more information, see [Navigator Encrypt key migration with HSM](#).

What's New in Schema Registry

Learn about the new features of Schema Registry in Cloudera Runtime 7.3.1.

Enable SMM principal as trusted proxy user in Schema Registry

SMM usually connects to Schema Registry on behalf of an end user. For requests coming from SMM, Schema Registry can now extract and authorize the end user to authorize the request.

New fingerprint version is added to Schema Registry with configuring option

A new fingerprint version, V2, is available in Schema Registry that contains the missing schema parts from the previous version. Newly created clusters use the V2 fingerprint version. Upgraded clusters still use the V1 fingerprint version, but the Fingerprint Version property can be used to change the fingerprint version in Schema Registry. Cloudera recommends changing the fingerprint version to V2 following an upgrade. For more information on how to change the fingerprint version, see [Step 9: Complete Post-Upgrade steps for upgrades to Cloudera Private Cloud Base](#).

This improvement fixes an issue in previous versions. For more information on the original issues as well as Schema Registry fingerprinting, see [TSB-713](#). Note that even if you switch to V2, some issues might still persist, see [TSB-718](#) for more information.

Support for additional JVM options

Additional JVM options can be passed to Schema Registry using the `schema.registry.additional.java.options` property in Cloudera Manager.

What's New in Apache Solr

Learn about the new features of Apache Solr in Cloudera Runtime 7.3.1.

Data Discovery and Exploration (Technical Preview)

A new Data Discovery and Exploration cluster definition is available in Data Hub. It lets you explore and discover data sets ad-hoc; doing relevance-based analytics over unstructured data (logs, images, text, PDFs, etc). The cluster definition deploys HDFS, Hue, Solr, Spark, Yarn, and ZooKeeper services. The cluster definition is available for AWS.

What's New in Apache Spark

Learn about the new features of Apache Spark in Cloudera Runtime 7.3.1.

Spark 2 removed from Cloudera Runtime

Spark 3 is the default Spark version in Cloudera Runtime. Spark 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.

Related Information

[Upgrading Spark](#)

[Migrating Spark Applications](#)

What's new in Streams Messaging Manager

Learn about the new features of Streams Messaging Manager (SMM) in Cloudera Runtime 7.3.1.

Validation for duplicate property keys in Kafka Connect connector configuration

When validating Kafka Connect connector configurations, a warning is displayed if the configuration contains duplicate property keys. Duplicate property keys are highlighted with orange. The form can still be validated with the warnings present, but if there are duplicates, you are notified that only the value of the last occurrence is used.

Search supports regular expressions

The search component on the Topics, Brokers, Consumers, Producers page can now perform a regexp search.

Visual clue when restarting on Kafka Connect

When clicking restart on Kafka Connect tasks or connectors, a loading circle is displayed in case of synchronous calls. The loading circle disappears once a response is received. For asynchronous calls, a pop-up is displayed, stating that the task or connector is restarted.

UX improvements

- Fixed text overflow in the side panel column headers
- Listing page table headers are now sticky of the nested table headers
- Listing page table styling has been improved for readability
- Filter selector drop-downs are now styled consistently
- Sidebar menu pop-ups are no longer hidden under tables
- Class names on the Kafka Connect popup are now wrapped into the containing pop-ups
- The password field is no longer obfuscated when using a file provider as a password
- Fixed the alignment of values on the Connector metrics page
- Source and sink connectors are now separate tabs on the connector creation modal
- Fixed visual issues on the topic creation modal
- Increased consistency in element contrast and text style throughout the UI
- Active and Inactive statuses now have high contrast
- The expand icon is now consistent throughout the UI

Expand security-related headers set by SMM

The following security related headers were added to SMM UI endpoints:

- Referrer-Policy
- Cross-Origin-Embedder-Policy

- Cross-Origin-Opener-Policy
- Cross-Origin-Resource-Policy

SMM uses trusted proxy authentication when connecting to Schema Registry

You can only interact with schemas through SMM if the necessary Ranger policies are set up for Schema Registry. For SMM UI, you must have the correct permissions to check messages deserialized with Avro on Data Explorer.

Dedicated endpoint for connector creation

A dedicated endpoint for connector creation is introduced. The endpoint is `POST /api/v1/admin/kafka-connect/connectors`. Requests made to this endpoint fail with the following error message if the connector name specified in the request exists.

```
{
  "error_code": 409,
  "message": "Connector [***NAME***] already exists"
}
```

In previous versions, `PUT /api/v1/admin/kafka-connect/connectors/{connector}` was the only endpoint you could use to create connectors. However, this endpoint is also used to update connectors. If you specify the name of an existing connector in the request, the endpoint updates existing connector. As a result, Cloudera recommends that you use the new endpoint for connector creation going forward. The SMM UI is also updated and uses the new endpoint for connector creation.

Jersey client timeout now configurable

SMM uses internal Jersey clients to make requests to Kafka Connect and Cruise Control. The connection and read timeouts for these clients was previously hard-coded to 30 seconds. Configuring them was not possible. This release introduces new properties, which enable you to configure the connection and read timeouts of these clients. The default timeout remains 30 seconds. The properties introduced are as follows.

- Kafka Connect Client Connect timeout
- Kafka Connect Client Read timeout
- Cruise Control Client Connect timeout
- Cruise Control Client Read timeout

What's New in Streams Replication Manager

Learn about the new features of Streams Replication Manager (SRM) in Cloudera Runtime 7.3.1.

The `--to` option in `srm-control` now creates the file if it does not exist

From now on, `srm-control` creates the file specified with the `--to` option if the file does not exist.

Configurations to customize replication-records-lag metric calculation

Three new properties are introduced that enable you to control how SRM calculates the replication-records-lag metric. This metric provides information regarding the replication lag based on offsets. The metric is available both on the cluster and the topic level. The following new properties are introduced because the calculation of the metric with default configurations might add latency to replications and impact SRM performance. While these properties are configured in Cloudera Manager, they do not have dedicated configuration entries. Instead, you add them to Streams Replication Manager's Replication Configs to configure them.

Property	Default Value	Description
<code>replication.records.lag.calc.enabled</code>	<code>true</code>	Controls whether the replication-records-lag metric is calculated. This metric provides information regarding the replication lag based on offsets. The metric is available both on the cluster and the topic level. The calculation of this metric might add latency to replications and impact SRM performance. If you are experiencing performance issues, you can try setting this property to <code>false</code> to disable the calculation of replication-records-lag. Alternatively, you can try fine-tuning how SRM calculates replication-records-lag with the <code>replication.records.lag.calc.period.ms</code> and <code>replication.records.lag.end.offset.timeout.ms</code> properties.
<code>replication.records.lag.calc.period.ms</code>	<code>0</code>	Controls how frequently SRM calculates the replication-records-lag metric. The default value of <code>0</code> means that the metric is calculated continuously. Cloudera recommends configuring this property to <code>15000 ms</code> (15 seconds) or higher if you are experiencing issues related to the calculation of replication-records-lag. A calculation frequency of 15 seconds or more results in the metric being available for consumption without any significant impact on SRM performance.
<code>replication.records.lag.end.offset.timeout.ms</code>	<code>6000</code>	Specifies the Kafka end offset timeout value used for calculating the replication-records-lag metric. Setting this property to a lower value than the default <code>6000 ms</code> (1 minute) might reduce latency in calculating replication-records-lag, however, replication-records-lag calculation might fail. A value higher than the default can help avoid metric calculation failures, but might increase replication latency and lower SRM performance.

What's New in Apache Hadoop YARN

Learn about the new features of Apache Hadoop Yarn in Cloudera Runtime 7.3.1.

Queue Manager

YARN Queue Manager is the queue management graphical user interface for Apache Hadoop YARN Capacity Scheduler. You can use YARN Queue Manager UI to manage your cluster capacity using queues to balance resource requirements of multiple applications from various users. Using YARN Queue Manager UI, you can set scheduler level properties and queue level properties. You can also view, sort, search, and filter queues using the YARN Queue Manager UI.

For more information about Queue Manager, see [Manage Queues](#).

FPGA as a resource type

You can use FPGA as a resource type. For more information, see [Use FPGA scheduling](#).

New configuration property to enable or disable the YARN recommendation engine APIs

The YARN Recommendation API now recommends scaling cluster nodes up or down based on the demand and idle state of cluster resources. This feature can be turned on/off using the YARN configuration property `yarn.cluster.scaling.recommendation.enable`.

Unaffected Components in this release

There are no new features for the following components in Cloudera Runtime 7.3.1.

- Avro
- Hadoop
- HBase
- HDFS
- Knox
- Kudu
- MapReduce
- Ozone
- Phoenix
- Parquet
- Sqoop
- Tez
- Zookeeper

What's new in Platform Support

You must be aware of the platform support for the Cloudera Runtime 7.3.1 release.

This section describes the platform support changes for the Cloudera Runtime 7.3.1 associated with Cloudera Private Cloud Base 7.3.1

Platform Support Enhancements

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

For more information on the support matrix, see [Cloudera Support Matrix](#).

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 should also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.11.1.7.3.1.0-197
Apache Atlas	2.1.0.7.3.1.0-197
Apache Calcite	1.25.0.7.3.1.0-197
Apache Avatica	1.22.0.7.3.1.0-197
Apache Avro	1.11.1.7.3.1.0-197
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.3.1.0-197
Apache HBase	2.4.17.7.3.1.0-197

Component	Version
Apache Hive	3.1.3000.7.3.1.0-197
Apache Iceberg	1.3.1.7.3.1.0-197
Apache Impala	4.0.0.7.3.1.0-197
Apache Kafka	3.4.1.7.3.1.0-197
Apache Knox	2.0.0.7.3.1.0-197
Apache Kudu	1.17.0.7.3.1.0-197
Apache Livy	0.7.23000.7.3.1.0-197
Apache MapReduce	3.1.1.7.3.1.0-197
Apache Ozone	1.3.0.7.3.1.0-197
Apache Oozie	5.1.0.7.3.1.0-197
Apache ORC	1.8.3.7.3.1.0-197
Apache Parquet	1.12.3.7.3.1.0-197
Apache Phoenix	5.1.3.7.3.1.0-197
Apache Ranger	2.4.0.7.3.1.0-197
Apache Solr	8.11.2.7.3.1.0-197
Apache Spark	3.4.1.7.3.1.0-197
Apache Sqoop	1.4.7.7.3.1.0-197
Apache Tez	0.9.1.7.3.1.0-197
Apache ZooKeeper	3.8.1.7.3.1.0-197

Other Components

Component	Version
Cruise Control	2.5.116.7.3.1.0-197
Data Analytics Studio	1.4.2.7.3.1.0-197
GCS Connector	2.1.2.7.3.1.0-197
Hue	4.5.0.7.3.1.0-197
Search	1.0.0.7.3.1.0-197
Schema Registry	0.10.0.7.3.1.0-197
Streams Messaging Manager	2.3.0.7.3.1.0-197
Streams Replication Manager	1.1.0.7.3.1.0-197
Data Discovery and Exploration	Technical Preview

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.3.1.0-197
Hive Meta Store (HMS)	1.0.0.7.3.1.0-197
Hive on Tez	1.0.0.7.3.1.0-197
Hive Warehouse Connector	1.0.0.7.3.1.0-197
Spark Atlas Connector	3.4.1.7.3.1.0-197
Spark Schema Registry	3.4.1.7.3.1.0-197



Note: Cloudera Ozone version 1.3.0 code is equivalent to Apache Ozone 1.4.0 in the 7.3.1 release. However, the version number will be reset in the next release.

Using the Cloudera Runtime Maven repository 7.3.1

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>
```

Runtime 7.3.1.0-197

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	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-classification-updater	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-client-common	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-client-v1	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-client-v2	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-common	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-distro	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-docs	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-graphdb-api	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-graphdb-common	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-graphdb-janus	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-hdfs-bridge	2.1.0.7.3.1.0-197

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-index-repair-tool	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-intg	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-notification	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-plugin-classloader	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-repository	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-server-api	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	atlas-testtools	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	hbase-bridge	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	hbase-bridge-shim	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	hbase-testing-util	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	hdfs-model	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	hive-bridge	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	hive-bridge-shim	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	impala-bridge	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	impala-bridge-shim	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	impala-hook-api	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	kafka-bridge	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	kafka-bridge-shim	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	navigator-to-atlas	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	sample-app	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	sqoop-bridge	2.1.0.7.3.1.0-197
Atlas	org.apache.atlas	sqoop-bridge-shim	2.1.0.7.3.1.0-197
Avro	org.apache.avro	avro	1.11.1.7.3.1.0-197
Avro	org.apache.avro	avro-android	1.11.1.7.3.1.0-197
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.3.1.0-197
Avro	org.apache.avro	avro-compiler	1.11.1.7.3.1.0-197
Avro	org.apache.avro	avro-grpc	1.11.1.7.3.1.0-197
Avro	org.apache.avro	avro-ipc	1.11.1.7.3.1.0-197
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.3.1.0-197
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.3.1.0-197
Avro	org.apache.avro	avro-mapred	1.11.1.7.3.1.0-197
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.3.1.0-197
Avro	org.apache.avro	avro-perf	1.11.1.7.3.1.0-197
Avro	org.apache.avro	avro-protobuf	1.11.1.7.3.1.0-197
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.3.1.0-197
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.3.1.0-197
Avro	org.apache.avro	avro-thrift	1.11.1.7.3.1.0-197
Avro	org.apache.avro	avro-tools	1.11.1.7.3.1.0-197

Project	groupId	artifactId	version
Avro	org.apache.avro	trevni-avro	1.11.1.7.3.1.0-197
Avro	org.apache.avro	trevni-core	1.11.1.7.3.1.0-197
Calcite	org.apache.calcite	calcite-babel	1.25.0.7.3.1.0-197
Calcite	org.apache.calcite	calcite-core	1.25.0.7.3.1.0-197
Calcite	org.apache.calcite	calcite-druid	1.25.0.7.3.1.0-197
Calcite	org.apache.calcite	calcite-linq4j	1.25.0.7.3.1.0-197
Calcite	org.apache.calcite	calcite-server	1.25.0.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-benchmark	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.3.1.0-197

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-minicluster	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.3.1.0-197

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.3.1.0-197
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.3.1.0-197
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-client	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-common	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-examples	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-http	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-it	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-logging	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-replication	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-rest	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-server	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-shell	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.3.1.0-197
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.3.1.0-197
Hive	org.apache.hive	catalogd-unit	3.1.3000.7.3.1.0-197

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-beeline	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-classification	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-cli	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-common	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-contrib	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-exec	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-iceberg-catalog	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-iceberg-handler	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-iceberg-shading	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-impala	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-it-iceberg	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-it-impala	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-it-util	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-jmh	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-metastore	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-parser	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-serde	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-service	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.3.1.0-197

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-shims	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-streaming	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-testutils	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-udf	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	kafka-handler	3.1.3000.7.3.1.0-197
Hive	org.apache.hive	patched-iceberg-api	patched-1.3.1.7.3.1.0-197-3.1.3000.7
Hive	org.apache.hive	patched-iceberg-core	patched-1.3.1.7.3.1.0-197-3.1.3000.7
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector-spark3_2.12	1.0.0.7.3.1.0-197
Kafka	org.apache.kafka	ci	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	connect	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	connect-api	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	connect-file	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	connect-json	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	generator	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.3.1.0-197

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.3.1.0-197
Kafka	org.apache.kafka	trogdor	3.4.1.7.3.1.0-197
Knox	org.apache.knox	gateway-adapter	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-admin-ui	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-applications	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-cloud-bindings	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-demo-ldap	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-demo-ldap-launcher	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-discovery-ambari	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-discovery-cm	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-docker	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-i18n	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-i18n-logging-log4j	2.0.0.7.3.1.0-197

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-i18n-logging-sl4j	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-openapi-ui	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-performance-test	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-ha	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-identity-assertion-common	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-identity-assertion-no-doas	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-jersey	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-rewrite	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-rewrite-common	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-security-authc-anon	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-security-authz-acls	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-security-authz-composite	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-security-authz-path-acls	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-security-clientcert	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-security-hadoopauth	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-security-jwt	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-security-pac4j	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-security-preauth	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-security-shiro	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-provider-security-webappsec	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-release	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-server	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-server-launcher	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-server-xforwarded-filter	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-admin	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-as	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-auth	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-definitions	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-hashicorp-vault	2.0.0.7.3.1.0-197

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-service-hbase	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-health	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-hive	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-idbroker	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-idbroker-plugins	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-impala	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-jkg	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-knoxsso	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-knoxssout	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-knoxtoken	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-livy	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-metadata	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-nifi	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-nifi-registry	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-remoteconfig	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-rm	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-session	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-storm	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-test	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-tgs	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-vault	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-service-webhdfs	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-shell	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-shell-launcher	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-shell-release	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-shell-samples	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-spi	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-spi-common	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-test	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-test-idbroker	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-test-release-utils	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-test-utils	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-topology-hadoop-xml	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-topology-simple	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-util-common	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-util-configinjector	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-util-launcher	2.0.0.7.3.1.0-197
Knox	org.apache.knox	gateway-util-urltemplate	2.0.0.7.3.1.0-197
Knox	org.apache.knox	hadoop-examples	2.0.0.7.3.1.0-197

Project	groupId	artifactId	version
Knox	org.apache.knox	knox-cli-launcher	2.0.0.7.3.1.0-197
Knox	org.apache.knox	knox-homepage-ui	2.0.0.7.3.1.0-197
Knox	org.apache.knox	knox-token-generation-ui	2.0.0.7.3.1.0-197
Knox	org.apache.knox	knox-token-management-ui	2.0.0.7.3.1.0-197
Knox	org.apache.knox	knox-webshell-ui	2.0.0.7.3.1.0-197
Knox	org.apache.knox	webhdfs-kerb-test	2.0.0.7.3.1.0-197
Knox	org.apache.knox	webhdfs-test	2.0.0.7.3.1.0-197
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.3.1.0-197
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.3.1.0-197
Kudu	org.apache.kudu	kudu-client	1.17.0.7.3.1.0-197
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.3.1.0-197
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.3.1.0-197
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.3.1.0-197
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.3.1.0-197
Livy	org.apache.livy	livy-api	0.7.23000.7.3.1.0-197
Livy	org.apache.livy	livy-client-common	0.7.23000.7.3.1.0-197
Livy	org.apache.livy	livy-client-http	0.7.23000.7.3.1.0-197
Livy	org.apache.livy	livy-core_2.12	0.7.23000.7.3.1.0-197
Livy	org.apache.livy	livy-examples	0.7.23000.7.3.1.0-197
Livy	org.apache.livy	livy-integration-test	0.7.23000.7.3.1.0-197
Livy	org.apache.livy	livy-repl_2.12	0.7.23000.7.3.1.0-197
Livy	org.apache.livy	livy-rsc	0.7.23000.7.3.1.0-197
Livy	org.apache.livy	livy-scala-api_2.12	0.7.23000.7.3.1.0-197
Livy	org.apache.livy	livy-server	0.7.23000.7.3.1.0-197
Livy	org.apache.livy	livy-test-lib	0.7.23000.7.3.1.0-197
Livy	org.apache.livy	livy-thriftserver	0.7.23000.7.3.1.0-197
Livy	org.apache.livy	livy-thriftserver-session	0.7.23000.7.3.1.0-197
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-analyzers-openslp	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.3.1.0-197

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-core	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-join	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.3.1.0-197
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.3.1.0-197
Oozie	org.apache.oozie	oozie-client	5.1.0.7.3.1.0-197
Oozie	org.apache.oozie	oozie-core	5.1.0.7.3.1.0-197
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.3.1.0-197
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.3.1.0-197
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.3.1.0-197
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.3.1.0-197
Oozie	org.apache.oozie	oozie-server	5.1.0.7.3.1.0-197
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.3.1.0-197
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.3.1.0-197
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.3.1.0-197
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.3.1.0-197
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.3.1.0-197
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.3.1.0-197
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.3.1.0-197
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.3.1.0-197
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.3.1.0-197
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.3.1.0-197
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.3.1.0-197
ORC	org.apache.orc	orc-core	1.8.3.7.3.1.0-197
ORC	org.apache.orc	orc-examples	1.8.3.7.3.1.0-197

Project	groupId	artifactId	version
ORC	org.apache.orc	orc-mapreduce	1.8.3.7.3.1.0-197
ORC	org.apache.orc	orc-shims	1.8.3.7.3.1.0-197
ORC	org.apache.orc	orc-tools	1.8.3.7.3.1.0-197
Ozone	org.apache.ozone	hdds-annotation-processing	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-client	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-common	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-config	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-container-service	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-crypto-api	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-crypto-default	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-docs	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-erasurecode	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-interface-admin	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-interface-client	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-interface-server	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-managed-rocksdb	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-rocks-native	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-server-framework	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-server-scm	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-test-utils	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	hdds-tools	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	mini-chaos-tests	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-client	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-common	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-csi	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-datanode	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-dist	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-filessystem	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-filessystem-common	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-filessystem-hadoop2	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-filessystem-hadoop3	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-filessystem-shaded	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-httpfsgateway	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-insight	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-integration-test	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-interface-client	1.3.0.7.3.1.0-197

Project	groupId	artifactId	version
Ozone	org.apache.ozone	ozone-interface-storage	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-manager	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-network-tests	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-recon	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-reconcodegen	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-s3-secret-store	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-s3gateway	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	ozone-tools	1.3.0.7.3.1.0-197
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.3.0.7.3.1.0-197
Parquet	org.apache.parquet	parquet-avro	1.12.3.7.3.1.0-197
Parquet	org.apache.parquet	parquet-cascading	1.12.3.7.3.1.0-197
Parquet	org.apache.parquet	parquet-cascading3	1.12.3.7.3.1.0-197
Parquet	org.apache.parquet	parquet-column	1.12.3.7.3.1.0-197
Parquet	org.apache.parquet	parquet-common	1.12.3.7.3.1.0-197
Parquet	org.apache.parquet	parquet-encoding	1.12.3.7.3.1.0-197
Parquet	org.apache.parquet	parquet-format-structures	1.12.3.7.3.1.0-197
Parquet	org.apache.parquet	parquet-generator	1.12.3.7.3.1.0-197
Parquet	org.apache.parquet	parquet-hadoop	1.12.3.7.3.1.0-197
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.12.3.7.3.1.0-197
Parquet	org.apache.parquet	parquet-jackson	1.12.3.7.3.1.0-197
Parquet	org.apache.parquet	parquet-pig	1.12.3.7.3.1.0-197
Parquet	org.apache.parquet	parquet-pig-bundle	1.12.3.7.3.1.0-197
Parquet	org.apache.parquet	parquet-protobuf	1.12.3.7.3.1.0-197
Parquet	org.apache.parquet	parquet-scala_2.12	1.12.3.7.3.1.0-197
Parquet	org.apache.parquet	parquet-thrift	1.12.3.7.3.1.0-197
Parquet	org.apache.parquet	parquet-tools	1.12.3.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.3.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.3.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix-core	5.1.3.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.1.6	5.1.3.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.2.5	5.1.3.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.3.0	5.1.3.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.0	5.1.3.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.1.3.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.0	5.1.3.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.4	5.1.3.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.3.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.3.1.0-197

Project	groupId	artifactId	version
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.3.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.3.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix5-spark3	6.0.0.7.3.1.0-197
Phoenix	org.apache.phoenix	phoenix5-spark3-shaded	6.0.0.7.3.1.0-197
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	jisql	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-gs-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.3.1.0-197

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-raz-hook-s3	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-raz-s3	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-raz-s3-lib	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-s3-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-trino-plugin	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-util	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.3.1.0-197

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	sample-client	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	sampleapp	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	shaded-raz-hook-s3	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.3.1.0-197
Ranger	org.apache.ranger	unixusersync	2.4.0.7.3.1.0-197
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.3.1.0-197
Solr	org.apache.solr	solr-analytics	8.11.2.7.3.1.0-197
Solr	org.apache.solr	solr-cell	8.11.2.7.3.1.0-197
Solr	org.apache.solr	solr-core	8.11.2.7.3.1.0-197
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.3.1.0-197
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.3.1.0-197
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.3.1.0-197
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.3.1.0-197
Solr	org.apache.solr	solr-langid	8.11.2.7.3.1.0-197
Solr	org.apache.solr	solr-ltr	8.11.2.7.3.1.0-197
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.3.1.0-197
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.3.1.0-197
Solr	org.apache.solr	solr-security-util	8.11.2.7.3.1.0-197
Solr	org.apache.solr	solr-solrj	8.11.2.7.3.1.0-197
Solr	org.apache.solr	solr-test-framework	8.11.2.7.3.1.0-197
Solr	org.apache.solr	solr-velocity	8.11.2.7.3.1.0-197
Spark	org.apache.spark	spark-avro_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-catalyst_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-connect-client-jvm_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-connect-common_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-connect_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-core_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-graphx_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-hadoop-cloud_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-hive_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-kubernetes_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-kvstore_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-launcher_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-mllib-local_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-mllib_2.12	3.4.1.7.3.1.0-197

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-network-common_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-network-shuffle_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-network-yarn_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-protobuf_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-repl_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-shaded-raz	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-sketch_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-sql-kafka-0-10_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-sql_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-streaming_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-tags_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-unsafe_2.12	3.4.1.7.3.1.0-197
Spark	org.apache.spark	spark-yarn_2.12	3.4.1.7.3.1.0-197
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.3.1.0-197
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.3.1.0-197
Tez	org.apache.tez	hadoop-shim	0.9.1.7.3.1.0-197
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.3.1.0-197
Tez	org.apache.tez	tez-api	0.9.1.7.3.1.0-197
Tez	org.apache.tez	tez-aux-services	0.9.1.7.3.1.0-197
Tez	org.apache.tez	tez-common	0.9.1.7.3.1.0-197
Tez	org.apache.tez	tez-dag	0.9.1.7.3.1.0-197
Tez	org.apache.tez	tez-examples	0.9.1.7.3.1.0-197
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.3.1.0-197
Tez	org.apache.tez	tez-history-parser	0.9.1.7.3.1.0-197
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.3.1.0-197
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.3.1.0-197
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.3.1.0-197
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.3.1.0-197
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.3.1.0-197
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.3.1.0-197
Tez	org.apache.tez	tez-tests	0.9.1.7.3.1.0-197
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.3.1.0-197
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.3.1.0-197
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.3.1.0-197
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.3.1.0-197
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.3.1.0-197

Project	groupId	artifactId	version
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.3.1.0-197
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.3.1.0-197
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.3.1.0-197
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.3.1.0-197
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.3.1.0-197
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.3.1.0-197
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.3.1.0-197
Zeppelin	org.apache.zeppelin	zeppelin-shaded-raz	0.8.2.7.3.1.0-197
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.3.1.0-197
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.3.1.0-197
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.3.1.0-197
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.3.1.0-197
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.3.1.0-197
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.3.1.0-197
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.3.1.0-197
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.3.1.0-197
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.3.1.0-197
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.3.1.0-197
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.3.1.0-197
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.3.1.0-197
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.3.1.0-197

Fixed Issues In Cloudera Runtime 7.3.1

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

Fixed Issues in Atlas

Review the list of Atlas issues that are resolved in Cloudera Runtime 7.3.1.

CDPD-69962: fetchType as "incremental" does full export instead of "CONNECTED"

Earlier, the first incremental export operation performed on a target entity used to fetch all entities even if they were not related to the targeted entity. This affected the performance as it imports more data than what was expected.

Now, the first incremental export will only fetch the entities which are related to the target entity. Also, if the target entity is connected to a lineage, then only the immediately connected entities in the lineage will get exported and not the whole lineage.

CDPD-67654: [Atlas] [navigator2atlas] Status of deleted table is ACTIVE in Atlas after navigator2atlas migration

Deleted hive tables migrated via the Navigator to Atlas transition may shown as active in Apache Atlas. Changes done in the Nav2Atlas module to set the relationType as hive_table_storagedesc of relationship attribute table for every entity of hive_storagedesc.

CDPD-72732: [UCL] Incorrect Atlas audits generated for updates with atlas.hook.hive.skip.dml.messages set to true/false in 7.3.0.1 CDP Private Cloud Base

The Apache Atlas property `atlas.hook.hive.skip.dml.messages = true` can be used to reduce the number of audits that are generated for any DML command executed over a hive entity.

The default value for `hive.split.update` is set to `true` in 7.3.1 causing two audits to be generated for one update command: one delete and one insert. This will impact Apache Atlas when `atlas.hook.hive.skip.dml.messages = false` (Atlas is processing Data Manipulation events) and `atlas.entity.audit.differential = false` (Atlas logs the full entity metadata during every update).

CDPD-71516: Temporarily disable the tasks tab on Entity Detail page

The Entity Detail page was showing "Something went wrong". This is occurring because on loading the Entity Detail page, an API call (`/api/atlas/admin/tasks`) is made to get all the tasks that are created when deferred actions features are enabled. The Entity Detail page task tab and task API will display in UI depending upon the server side property `atlas.tasks.ui.tab.enabled`. Initially, this is set to `false`. Therefore, temporarily the task tab on entity detail page in UI is disabled.

Apache Jira: [ATLAS-4880](#)

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

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

OPSAPS-68461: Update GC and JVM options for Atlas service for supporting JDK17 in main Atlas CSD

The issue of existing ATLAS OPTS not working for JDK17 is fixed.

Fixed Issues in Apache Avro

There are no fixed issues for Apache Avro in Cloudera Runtime 7.3.1.

Apache patch information

None

Fixed issues in Cloud Connectors

Review the list of Cloud Connectors issues that are resolved in Cloudera Runtime 7.3.1.

Apache patch information

- HADOOP-18855 - Tuning and stabilization of Vector IO
- MAPREDUCE-7474 - Improve commit resilience and performance in Manifest Committer for ABFS

Fixed issues in Cruise Control

There are no fixed issues for Cruise Control in Cloudera Runtime 7.3.1.

Fixed Issues in Hadoop

There are no fixed issues for Hadoop in Cloudera Runtime 7.3.1.

Fixed Issues in HDFS

Review the list of HDFS issues that are resolved in Cloudera Runtime 7.3.1.

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

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

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

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

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

This issue is now resolved.

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

This issue is now resolved by fixinf service.sdl.

CDPD-67823: Ranger RMS gives all permissions to the user through the Create permission.

This issue is now resolved. An additional check is added to ensure that the user attempting to alter any HDFS directory that maps to the Hive database is the owner of the Hive database.

Fixed Issues in HBase

Review the list of HBase issues that are resolved in Cloudera Runtime 7.3.1.

CDPD-67520: JWT authentication expects [sub] claim in the payload

A JWT payload can have a custom claim for Subject/Principal instead of the standard sub claim.

You can set the hbase.security.oauth.jwt.token.principal.claim configuration property in Cloudera Manager under HBase Service Advanced Configuration Snippet (Safety Valve) for hbase-site.xml to define the custom Subject/Principal claim.

CDPD-66387: RegionServer should be aborted when WAL.sync throws TimeoutIOException

This fix adds additional logic for WAL.sync. If WAL.sync gets a timeout exception, HBase wraps TimeoutIOException as a special WALSyncTimeoutIOException. When the upper layer such as HRegion.doMiniBatchMutate called by HRegion.batchMutation catches this special exception, HBase aborts the region server.

Apache Jira: [HBASE-27230](#)

CDPD-65373: Make delay prefetch property dynamically configurable

This change allows you to dynamically configure the hbase.hfile.prefetch.delay property using the Cloudera Manager. You need to update the value and refresh the HBase service. The new value is applied to the HBase service automatically.

Apache Jira: [HBASE-28292](#)

CDPD-74494: JVM crashes intermittently on ARM64 machines

After noticing the JVM crashes in the HBase service that is based on arm64 architecture and uses JDK 17, the fix is applied that refactors the module and the large implementation function into multiple smaller functions. The issue was observed in a specific module that had a very large member function.

Apache Jira: [HBASE-28206](#)

CDPD-73117: Bucket cache utilization is dropped after a rolling restart

For a persistent bucket cache of a size higher than 1.3 TB, the corresponding backing-map information (information related to the persistence cache) grows beyond 2 GB. But, 2 GB is the limit of the protobuf message sizes. These protobuf messages are used to persist the backing map information. If the size of the message grew beyond 2 GB, the backing map partially persisted and after a restart, the size of the cache seemed to be reduced.

With this fix, backing map information was chunked in smaller chunks with sizes below 2 GB. Now all information, even beyond 2 GB, is persisted and can be retrieved back after a rolling restart.

OPSAPS-70946: The hbase-site.xml file does not contain xinclude for the refreshable files

HBase supports generating hbase-site.xml with xinclude which is needed for the hbase-site-refreshable.xml file.

OPSAPS-70908: Refresh cluster command fails during ephemeral cache zero downtime upgrade

Configurations from refreshable files encountered authentication failure during the refresh command when Kerberos is enabled.

```
hbase/hbase.sh
[ "refresh-regionserver", "hbase.hfile.prefetch.delay", "hbase.rs.cacheblocksonwrite",
  "hbase.block.data.cacheonread", "hbase.rs.evictblocksonclose" ]
```

To fix this, RegionServerRefreshCommand now sets SCM_KERBEROS_PRINCIPAL as the Kerberos principal in the region server refresh process in the environment.

OPSAPS-70866: Invalid HBase prefetch configurations during rolling runtime upgrade

The default values of hbase_hfile_prefetch_delay and hbase.block.data.cacheonread are reverted to 1000 ms and are set to true.

OPSAPS-70294: HBase must use load balancing for the WEBHBASE Knox service

For CDPD 7.3.0 and later, the WEBHBASE service is configured for sticky load balancing instead of high availability in Knox.

OPSAPS-70035: HBase ZooKeeper client TLS toggle should also control the daemon roles

This issue is fixed. HBase ZooKeeper secure client mode now affects all roles.

OPSAPS-69983: Set Zookeeper store types to HBase service configuration

HBase now automatically sets the ZooKeeper truststore type based on ScmParams.

OPSAPS-69805: HBase client configuration does not use a secure port if Client TLS is enabled

HBase only uses a secure ZooKeeper port in client connections if enabled explicitly.

OPSAPS-69757: Make HBase TLS connection to ZooKeeper disabled by default

The HBase TLS connection to ZooKeeper must be disabled because it breaks some use cases. Instead, HBase introduces a new property to enable or disable in client roles. The default value is disabled.

OPSAPS-57937: No alerts are generated when the HBase process is in a hung state

HBase master monitoring (canary) showed green status even if the master has not initialized yet and added extra checks to query HBase if it is up and running.

OPSAPS-53851: ZooKeeper SSL/TLS support for HBase

Cloudera Manager configures HBase for a secure ZooKeeper connection if ZooKeeper TLS is enabled.

CDPD-74725: HBase throws org.apache.hbase.thirdparty.io.netty.util.ResourceLeakDetector exception

HBase direct memory buffer leak issues are fixed which could lead to heap issues in the long run.

Apache Jiras: [HBASE-28890](#) and [HBASE-28893](#)

CDPD-72120: Allow specifying a filter for the REST multiget endpoint (addendum: add back SCAN_FILTER constant)

HBase allows specifying a filter for the REST multiget endpoint (addendum: add back SCAN_FILTER constant).

Apache Jira: [HBASE-28518](#)

CDPD-71008: REST Java client library assumes stateless servers

This issue is fixed.

Apache Jira: [HBASE-28500](#)

CDPD-71007: hbase-rest client shading conflicts with hbase-shaded-client in HBase 2.x

This issue is fixed.

Apache Jira: [HBASE-28526](#)

CDPD-71006: Support non-SPNEGO authentication methods and implement session handling in the REST Java client library

This issue is fixed.

Apache Jira: [HBASE-28501](#)

CDPD-70493: MultiRowRangeFilter deserialization fails in org.apache.hadoop.hbase.rest.model.ScannerModel

This issue is fixed.

Apache Jira: [HBASE-28626](#)

CDPD-69335: Use a single GET call in the REST multiget endpoint

This issue is fixed.

Apache Jira: [HBASE-28523](#)

CDPD-68900: HBase properties need to be dynamically configured

The following configurations can be dynamically configured.

- hbase.rs.evictblocksonclose
- hbase.rs.cacheblocksonwrite
- hbase.block.data.cacheonread

After changing values of these configurations restarting region servers is no longer required. These configurations help in getting better throughput.

Newly changed values in the hbase-site.xml are read by HBase and values in appropriate classes are updated.

CDPD-68550: BucketCache.notifyFileCachingCompleted might incorrectly consider a file fully cached

This issue is fixed.

Apache Jira: [HBASE-28458](#)

CDPD-68154: BuckeCache.evictBlocksByHfileName does not work after a cache recovery from a file

This issue is fixed.

Apache Jira: [HBASE-28450](#)

CDPD-64046: BucketCache.blocksByHFile might leak on allocationFailure or if encountering input/output errors can lead to cache leak and extra heap usage

This issue is fixed.

Apache Jira: [HBASE-28211](#)

CDPD-63765: Move the NavigableSet add operation to the writer thread in BucketCache

This issue fixes potential cache leaks and extra memory usage.

Apache Jira: [HBASE-26305](#)

CDPD-62737: PrefetchExecutor must not run for files from the CF levels that have disabled BLOCKCACHE

This fix allows disabling the caching or pre-caching of individual tables.

Apache Jira: [HBASE-28217](#)

CDPD-45890: Fix the miss count in one of the CombinedBlockCache getBlock implementations

This fix impacts the hit ratio chart's accuracy in Cloudera Manager.

Apache Jira: [HBASE-28189](#)

Fixed Issues in Hive

Review the list of Hive issues that are resolved in Cloudera Runtime 7.3.1.

CDPD-13406: Disable TopN in ReduceSinkOp when TopNKey is introduced

When both the ReduceSink and TopNKey operators are used together in a query, they both perform Top-N key filtering. This results in the same filtering logic being applied twice, causing slower query execution in query execution.

The Top-N key filtering logic within the ReduceSink operator is now disabled when the TopN Key operator is introduced. The patch ensures that only the TopNKey operator handles the Top-N filtering, while the other functionalities of the ReduceSink operator remain unaffected.

Apache Jira: [HIVE-23736](#)

CDPD-28339: Skip extra work in Cleaner when queue is empty

The Cleaner previously made unnecessary database calls and logged activities even when there were no candidates for cleaning.

This was optimized by skipping the extra DB calls and logging when the cleaning queue is empty, improving performance.

Apache Jira: [HIVE-24754](#)

CDPD-28174: Compaction task reattempt fails due to FileAlreadyExistsException

The issue arises when a compaction task is relaunched after the first attempt fails, leaving behind temporary directories. The second attempt encounters a FileAlreadyExistsException because the `_tmp` directory created during the first attempt was not cleared.

The solution ensures that compaction reattempts clear the old files from previous attempts before starting, preventing the failure caused by stale directories.

Apache Jira: [HIVE-24882](#), [HIVE-23058](#)

CDPD-45285: Incorrect results for IN UDF on Parquet columns of CHAR/VARCHAR type

Queries with case statements and multiple conditions return incorrect results for tables in Parquet format, particularly with CHAR/VARCHAR types. The issue is not observed with ORC or TextFile formats and can be bypassed by setting `hive.optimize.point.lookup` to false.

The issue was addressed by adding the necessary CASTs during IN clause conversion.

Apache Jira: [HIVE-26320](#)

CDPD-24412: Compaction queue entries stuck in 'ready for cleaning' state

When multiple compaction tasks run simultaneously on the same table, only one task removes obsolete files while others remain in the ready for cleaning state, leading to an accumulation of queue entries.

Add a mechanism to automatically clear or re-evaluate entries stuck in the ready for cleaning state to improve compaction task efficiency.

Apache Jira: [HIVE-25115](#)

CDPD-27291: getCrossReference fails when retrieving constraints from the primary key side

When retrieving constraints from the primary key side, the foreign key is passed as null, causing the operation to fail with a Db name cannot be null exception, especially when the metadata cache is enabled by default.

This has been resolved by ensuring that the foreign key constraint is correctly handled even when passed as null during constraint retrieval from the primary key side.

CDPD-15269: Add caching support for frequently called constraint APIs in catalogd's HMS interface

The `get_unique_constraints`, `get_primary_keys`, `get_foreign_keys`, and `get_not_null_constraints` APIs are called frequently during query compilation, particularly with TPCDS queries. Without caching, this leads to performance overhead.

Introduced caching for the above APIs in the Catalogd's HMS interface by adding `ValidWriteIdList` and `tableId` to the API requests. This ensures that the cache or backing DB is appropriately used to serve responses.

Apache Jira: [HIVE-23931](#)

DWX-8663: ShuffleScheduler should report the original exception when shuffle becomes unhealthy

The `ShuffleScheduler` does not report the original exception when the shuffle becomes unhealthy, making it harder to diagnose the underlying issue.

This issue is now fixed.

Apache Jira: [TEZ-4342](#)

CDPD-43837: MSSQL upgrade scripts fail when adding TYPE column to DBS table

Schema upgrade for MSSQL fails with an error when trying to add the `TYPE` column to the DBS table due to the incorrect usage of the keyword `NATIVE` in the default value.

The issue was addressed by modifying the schema upgrade script to use a valid constant expression for the default value in MSSQL.

Apache Jira: [HIVE-25551](#)

CDPD-43890: Drop data connector if not exists should not throw an exception

The `DROP DATA CONNECTOR IF NOT EXISTS` command incorrectly throws a `NoSuchObjectException` when the connector does not exist.

The issue was addressed by ensuring that no exception is thrown if the `ifNotExists` flag is true during the drop operation.

Apache Jira: [HIVE-26299](#)

CDPD-43838: Filter out results for show connectors in Hive Metastore client side

The `SHOW CONNECTORS` command does not filter results based on authorization, such as Ranger policies, on the client side.

The issue was addressed by implementing client-side filtering in HMS to ensure that only connectors authorized by policies like Ranger are displayed.

Apache Jira: [HIVE-26246](#)

CDPD-43952: HMS get_all_tables method does not retrieve tables from remote database

The `get_all_tables` method in Hive Metastore handler only retrieves tables from the native database, unlike the `get_tables` method, which can retrieve tables from both native and remote databases.

The issue was addressed by updating the `get_all_tables` method to retrieve tables from both native and remote databases, ensuring consistency with the `get_tables` method.

Apache Jira: [HIVE-26171](#)

CDPD-55914: Select query on table with remote database returns NULL values with PostgreSQL and Redshift data connectors

Few datatypes are not mapped from Postgres or Redshift to Hive data types in the connector, which resulted in displaying null values for the columns of those data types.

This issue is fixed.

Apache Jira: [HIVE-27316](#)

CDPD-31726: Prevent NullPointerException by Checking Collations Return Value

A NullPointerException occurs during execution of an EXPLAIN cbo on a subquery when using Tez as the execution engine, leading to empty explain output.

Added a check for null return values from RelMetadataQuery.collations() to prevent NullPointerExceptions in RelFieldTrimmer and HiveJoin, ensuring stability during query execution.

Apache Jira: [HIVE-25749](#)

CDPD-27418: Incorrect row order after query-based MINOR compaction

The query-based MINOR compaction used an incorrect sorting order, which led to duplicated rows after multiple merge statements.

The sorting order was corrected, ensuring proper row handling.

Apache Jira: [HIVE-25258](#)

CDPD-27419: Incorrect row order validation for query-based major compaction

The row order validation for query-based MAJOR compaction incorrectly checked the order as bucketProperty, leading to failures with multiple bucketProperties.

The validation was updated to correctly check the order as originalTransactionId, bucketProperty, and rowId, and an improved error message was implemented.

```
Error: org.apache.hadoop.hive.ql.metadata.HiveException: Wrong sort order of Acid rows detected for the rows
```

Apache Jira: [HIVE-25257](#)

Enable proper handling of non-default schemas in Hive for JDBC databases

Hive fails to create an external table for a JDBC database when the table is in a non-default schema, causing PSQLErrorException error that the table does not exist.

Improved handling of tables in non-default schemas by correctly using the hive.sql.schema property. This ensures the table is found, preventing the error.

Apache Jira: [HIVE-25591](#)

CBO failure when using JDBC table with password through dbcp.password.uri

When a table is created using JDBCStorageHandler and the JDBC_PASSWORD_URI is specified for the password, the Cost-Based Optimizer (CBO) fails. This causes all the data to be fetched directly from the database and processed in Hive, impacting performance.

Adjustments were made to ensure CBO functions correctly when JDBC_PASSWORD_URI is used, allowing for proper optimization and preventing unnecessary data fetch from the database.

Apache Jira: [HIVE-25626](#)

CDPD-28904: Intermittent Hive JDBC SSO failures in virtual environments

Browser-based SSO with the Hive JDBC driver fails in virtual environments (like Windows VMs). The driver sometimes misses POST requests with the SAML token due to a race condition, causing authentication failures.

Resolved a race condition in the JDBC driver to ensure it properly handles SSO authentication in virtual environments, preventing POST request failures.

Apache Jira: [HIVE-25479](#)

CDPD-43672: Remove unnecessary optimizations in canHandleQbForCbo

The `canHandleQbForCbo()` includes an optimization where it returns an empty string if INFO logging is disabled, which complicates the logic and doesn't significantly impact performance.

The issue was addressed by simplifying the code in `canHandleQbForCbo()` and removing the unnecessary optimization related to logging.

Apache Jira: [HIVE-26438](#)

DWX-7648: Infinite loop during CBO parsing cause OOM in HiveServer2

HiveServer became unstable due to an infinite loop during query parsing with UNION operations, causing an out-of-memory error during cost-based and logical optimization phase. The issue occurred because Hive's custom metadata provider was not initialized.

The initialization has now been moved before CBO requires it.

Apache Jira: [HIVE-25220](#)

CDPD-31200: Reader not closed after check in AcidUtils, leading to resource exhaustion

The Reader in `AcidUtils.isRawFormatFile` is not being closed after the check is finished. This causes issues when resources on the DFSCient are limited, leading to connection pool timeouts such as Timeout waiting for connection from pool.

The fix includes automatically closing the Reader in `AcidUtils.isRawFormatFile`, which ensures that resources are freed up and prevents connection pool timeout issues.

Apache Jira: [HIVE-25683](#)

DWX-10336: SSL certificate import error in HiveServer2 with JWT authentication

JWT support for HiveServer, SSL certificate import fails due to self-signed certificates not being accepted by the JVM in environments. The error occurs during the initialization of the HTTP server.

The fix includes introducing a property to disable SSL certificate verification for downloading JWKS (JSON Web Key Set) in environments. This helps users bypass certificate validation.

Apache Jira: [HIVE-26425](#)

CDPD-42686: Query-based compaction fails for tables with complex data types and reserved keywords

Query-based compaction fails on tables with complex data types and columns containing reserved keywords due to incorrect quoting of column names when creating a temporary table.

The issue was addressed by ensuring that columns with reserved keywords are correctly quoted during the creation of temporary tables.

Apache Jira: [HIVE-26374](#)

CDPD-54605: HiveSchemaTool to honor metastore-site.xml for initializing metastore schema

The HiveSchemaTool fails to recognize the `metastore-site.xml` configuration when initializing the metastore schema. It defaults to using an embedded database instead of the specified MySQL database.

The issue was addressed by updating the HiveSchemaTool to ensure it properly reads the `metastore-site.xml` file, allowing for correct initialization of the metastore schema with the intended database configuration.

Apache Jira: [HIVE-26402](#)

CDPD-55135: JDBC data connector queries to avoid exceptions at CBO stage

JDBC data connector queries throw exceptions at the CBO stage due to incorrect handling of database, schema and table names. When querying, the database name is improperly swapped with the schema name, leading to error:

```
schema dev does not exist
```

The issue was addressed by changing the `hive.sql.table` property value from `databasename.tablename` to `tablename` and adding `hive.sql.table` property with `databasename`. This adjustment ensures that the CBO stage retrieves JDBC table information correctly, eliminating the errors related to schema and table name resolution.

Apache Jira: [HIVE-26192](#)

DWX-8296: Hive query vertex failure due to Kerberos authentication error

Hive queries fail during LLAP execution with vertex failures. The query-executor fails to communicate with the query-coordinator due to an authentication error in Kerberos.

Address the Kerberos authentication failure between query-executors and the query-coordinator to ensure proper task execution and prevent vertex failures during LLAP execution.

CDPD-45607: Atomic schema upgrades for HMS to prevent partial commits

SchemaTool may leave the metastore in an invalid state during schema upgrades because each change is autocommitted. If an upgrade fails mid-process, the schema is left partially updated, causing issues on reruns.

The issue was addressed by ensuring schema changes are committed only after the entire upgrade process completes successfully. If any step fails, no changes are applied, preventing partial updates and keeping the schema intact.

Apache Jira: [HIVE-25707](#)

CDPD-49232: Auto-reconnect data connectors after timeout

When data connectors remain idle for long, the JDBC connection times out. This requires a restart to re-establish the connection, rendering the connector unusable until then.

The issue was addressed by automatically checking if a connection is closed and re-establishing it when necessary. This ensures the connectors stay functional without needing a restart, and includes setting connection timeout and retry properties for more reliable reconnections.

Apache Jira: [HIVE-26045](#)

CDPD-49494: Allow JWT and LDAP authentication to co-exist in HiveServer2 configuration

Setting `hive.server2.authentication=JWT,LDAP` fails with a validation error, preventing HiveServer2 from starting due to conflicts between authentication types.

The issue was addressed by updating the validation logic to support JWT authentication alongside LDAP, ensuring HiveServer2 can start with both auth mechanisms enabled.

Apache Jira: [HIVE-26045](#)

CDPD-40732: Timestamps when reading parquet files with Vectorized reader

Timestamp shifts occur when reading Parquet files that were created in older Hive versions and vectorized execution is enabled. The vectorized reader is not able to exploit the metadata inside the Parquet file to apply the correct conversion. For instance, a timestamp written as 1880-01-01 00:00:00 may be read as 1879-12-31 23:52:58; the exact shift depends on the JVM timezone. The non-vectorized reader is not affected.

The fix ensures both vectorized and non-vectorized readers use the same logic to determine the correct timestamp conversion based on metadata and configuration.

Apache Jira: [HIVE-26270](#)

Fixed Issues in Hue

Review the list of Hue issues that are resolved in Cloudera Runtime 7.3.1

CDPD-65034: Receiving Error " TypeError: 'NoneType' object is not callable" in TCLIService.py when custom headers are being set

When XSRF (Cross-Site Request Forgery) or CSRF (Cross-Site Request Forgery) is enabled in Hive or Impala, you might encounter the error "Error " TypeError: 'NoneType' object is not callable" in TCLIService.py. You can resolve this issue by upgrading to 7.1.9 SP1 CHF3 or 7.3.1 versions.

Fixed Issues in Impala

Review the list of Impala issues that are resolved in Cloudera Runtime 7.3.1.

CDPD-69856: SIGSEGV crash while accessing query state from concurrent access during query execution

A crash can occur due to concurrent updates and reads of execution state, such as through the WebUI, during query processing.

Ensured atomic updates of execution state to prevent conflicts and crashes during concurrent operations.

Apache Jira: [IMPALA-12747](#)

Incorrect length when inserting multiple CHAR(N) values

When all values in a VALUES clause column are of type CHAR, but have differing lengths, Impala selects a common CHAR(max(lengths)) type, padding shorter values with spaces. This padding can lead to unexpected results if the destination column is VARCHAR or STRING, as inserting the values individually would not produce the same padded output.

The VALUES_STMT_AVOID_LOSSY_CHAR_PADDING query option was introduced to address this discrepancy. When set to true, this option prevents padding by automatically casting values to VARCHAR with the length of the longest value in the column, provided all values are CHAR and not of equal length. By default, this option is set to false.

Apache Jira: [IMPALA-10753](#)

Type mismatch in set operations with ALLOW_UNSAFE_CASTS enabled

The fix ensures that string literals are converted to numeric types correctly by considering the target type when ALLOW_UNSAFE_CASTS is enabled.

Apache Jira: [IMPALA-12285](#)

Fixed Issues in Apache Iceberg

Review the list of iceberg issues that are resolved in Cloudera Runtime 7.3.1.

CDPD-49011: Support incremental materialized view rebuild

This fix provides support for incremental materialized view rebuild when the Iceberg source tables only have the INSERT operation.

Apache JIRA: [HIVE-27101](#)

CDPD-48395: Upgrade the Parquet version to 1.12.3 for Hive

This fix upgrades the Parquet version for Hive to 1.12.3, which is the same Parquet version that is used for Iceberg.

Fixed Issues in Kafka

Review the list of Kafka issues that are resolved in Cloudera Runtime 7.3.1.

CDPD-65649: ReplicaAlterLogDirs stuck with Offset mismatch for the future replica

This is a backported fix, see [KAFKA-9087](#) for more information.

CDPD-66986: Mirrormaker 2 auto.offset.reset=latest not working

This is a backported fix, see [KAFKA-13988](#) for more information.

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

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

CDPD-71433: Connect logical type null values are not handled in AvroConnectTranslator

When the `time.precision.mode` property is set to `connect` for the Debezium connector, the connect logical types are used and null values are now handled.

CDPD-62059: AvroConnectTranslator should handle null values in fromConnectData method

Fix possible NPE exception issues in connector tasks which operate with Avro data format using the `Cloudera AvroConverter`.

Fixed Issues in Kudu

Review the list of Kudu issues that are resolved in Cloudera Runtime 7.3.1

KUDU-3576: Fix the Connection timeout After Tablet Server Restart

In Kudu, if a Java client application maintains an open connection to a tablet server and the tablet server is restarted or encounters a network error, the client cannot re-establish communication with the tablet server even after it comes back online. The fix resolves the issue by updating the Kudu Java client.

Apache Jira: [KUDU-3576](#)

KUDU-3577: Fix the altering tables with custom hash schemas

This fix resolves the issue caused by the design defect related to partition boundaries for ranges with custom hash schemas, which requires re-encoding the information as a part of `PartitionSchemaPB` stored in the system catalog upon particular modifications of the table's schema. A proper solution would be using the primary-key-only projection of a table's schema to encode the information on range boundaries while also maintaining backward compatibility with already-released Kudu clients.

Apache Jira: [KUDU-3577](#)

KUDU-3496: Support for SPNEGO dedicated keytab

Kudu now supports configuring a dedicated Spnego keytab.

Apache Jira: [KUDU-3496](#)

KUDU-3497: Optimize leader rebalancer algorithm

Optimized the leader balancing algorithm to effectively handle corner cases detailed in the Jira.

Apache Jira: [KUDU-3497](#)

KUDU-3486: Clearing Tombstoned Replica Metadata During Heartbeat Processing in Kudu TServer

The metadata of tombstoned replicas can occupy ``kudu-tserver`` memory and is now cleared during heartbeat processing.

Apache Jira: [KUDU-3486](#)

Fixed Issues in Apache Knox

Review the list of Knox issues that are resolved in CDP Private Cloud Base 7.3.1.

CDPD-73275: HTTP 404 responses while Knox is redeploying topologies

While you were redeploying topologies, Knox returned HTTP 404 responses.

Knox no longer returns HTTP 404 responses during topology redeployment, but returns HTTP 503 instead.

CDPD-70313: KNOX did not send Authentication header on FIPS configuration

KNOX neither sent the authentication header nor hadoop.auth cookie that was why the SMM UI sent back the HTTP 401 response and set the "www-authenticate": "Negotiate" header. Because of this, the SMM UI was inaccessible through Knox.

This issue is fixed now.

CDPD-67478: Custom topologies cannot be deleted

You cannot delete custom topologies that were created.

This issue is fixed.

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

This issue is fixed now.

CDPD-69305: /plugins/policies/importPoliciesFromFile API returns 500 service connectivity error through Knox Proxy

The fix imports large policy files using the Ranger importPoliciesFromFile API through Knox.

Apache patch information

- KNOX-3073
- KNOX-3058
- KNOX-3055
- KNOX-3054
- KNOX-3053
- KNOX-3052
- KNOX-3050
- KNOX-3049
- KNOX-3045
- KNOX-3040
- KNOX-3038
- KNOX-3037
- KNOX-3036
- KNOX-3029
- KNOX-3028
- KNOX-3026
- KNOX-3024
- KNOX-3023
- KNOX-3022
- KNOX-3020
- KNOX-3019
- KNOX-3018
- KNOX-3017
- KNOX-3016
- KNOX-3012
- KNOX-3007
- KNOX-3006
- KNOX-3005
- KNOX-3002

- KNOX-3001
- KNOX-3000
- KNOX-2994
- KNOX-2985
- KNOX-2983
- KNOX-2980
- KNOX-2979
- KNOX-2978
- KNOX-2976
- KNOX-2975
- KNOX-2974
- KNOX-2973
- KNOX-2972
- KNOX-2971
- KNOX-2970
- KNOX-2969
- KNOX-2968
- KNOX-2966
- KNOX-2963
- KNOX-2961
- KNOX-2960
- KNOX-2959
- KNOX-2958
- KNOX-2955
- KNOX-2951
- KNOX-2949
- KNOX-2948
- KNOX-2947
- KNOX-2946
- KNOX-2929
- KNOX-2896
- KNOX-2881

Fixed Issues in Apache Livy

Review the list of Livy issues that are resolved in Cloudera Runtime 7.3.1.

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

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

CDPD-73324 - LIVY_FOR_SPARK3 goes into down with Invalid Keystore format error in FIPS cluster

Fixed an issue that caused LIVY_FOR_SPARK3 to go into a bad state with Invalid Keystore format error in a 7.3.1 FIPS cluster.

Apache patch information

None

Fixed Issues in Navigator Encrypt

Review the list of Navigator Encrypt issues that are resolved in CDP Private Cloud Base 7.3.1.

CDPD-70115: Ranger KMS with Oracle DB was not supported for Navigator Encrypt

Navigator Encrypt deposit registration was failing with Ranger KMS DB with Oracle DB setup with the following error:

```
java.sql.SQLException: ORA-02289: sequence does not exist
Error Code: 2289
```

The issue is fixed now.

Fixed Issues in Oozie

Review the list of Oozie issues that are resolved in Cloudera Runtime 7.3.1.

CDPD-70422: Cannot enforce Oozie parameter oozie.http.hostname

A new property named oozie.http.hostname.override is now introduced to specify the interface that the Oozie Server must be using.

CDPD-71117: Oozie server does not pass action start time to action conf causes a restarting launcher doesn't find child apps

Whenever Yarn restarted the Oozie Launcher AM, Oozie could not find the previously started child jobs due to a missing original start timestamp from the Oozie Server. And the previously started child Jobs were not terminated when the Launcher AM was restarted. This issue is now resolved.

CDPD-48664: Retry mechanism anomaly in Oozie with High Availability enabled

There was an issue with the retry mechanism in Oozie when High Availability was enabled. This issue is now resolved.

CDPD-49745: Expand app_path column in *_JOBS tables to allow HDFS paths longer than 255 characters

The APP_PATH column now supports storing paths longer than 255 characters.

Fixed issues in Ozone

Review the list of Ozone issues that are resolved in Cloudera Runtime 7.3.1.

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

This issue is now resolved.

OPSAPS-70288: Improvements in master node decommissioning.

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

CDPD-74756: Update Ratis to 3.1.1

Updated Ratis dependency version from 3.1.0 to 3.1.1.

Apache Jira: [HDDS-11504](#)

CDPD-74241: OmSnapshotPurge should be in a different Ozone manager double buffer batch.

This issue is now resolved.

Apache Jira: [HDDS-11453](#)

CDPD-74200: Recon UI shows incorrect data about volume, bucket, and keys. Recon is unable to sync its data with OM DB.

This issue is now resolved.

CDPD-74074: The /v1/triggerdbsync/om api is working with non-admin user even if security is enabled.

This issue is now resolved.

Apache Jira: [HDDS-11436](#)

CDPD-73775: Replace solr.version with solr_lkgb_jar_version for Ozone to use downstream version of Solr.

Replaced pom solr.version with solr_lkgb_jar_version for Ozone to use downstream version of Solr.

CDPD-73447: Incorrect number of deleted containers shown in Recon UI.

The addition of the EMPTY_MISSING state to the retainOrUpdateRecord method enables Recon to correctly identify and manage the lifecycle of these containers, ensuring that no stale or deleted containers remain in memory or in Recon's records.

CDPD-73330: The namespace quota and namespace dist commands fail and displays the Path not found in the system error for the existing volume or bucket.

Fixed the Ozone admin namespace summary.

Apache Jira: [HDDS-10581](#)

CDPD-72142: Keys from DeletedTable and DeletedDirTable of Active Object Store (AOS) should be deleted on batch operation while creating a Snapshot.

On snapshot creation, DeletedTable and DeletedDirTable of AOS is cleared. This operation is not performed in the same transaction as Snapshot create which can cause an orphan block objects in case of bootstrapping and lagging follower. This issue is now resolved and Snapshot creation and clearing of the DeletedTable and DeletedDirTable on AOS will be a single batch operation.

Apache Jira: [HDDS-11183](#)

CDPD-72076: The OMDoubleBuffer error is displayed when handling OMRequest: cmdType: SnapshotMoveDeletedKeys.

This fixes OM crash issue when the follower is lagging and it executes purgeKeys or snapshotMoveDeletedKeys for the Snapshot in the one transaction.

Apache Jira: [HDDS-11152](#)

CDPD-72019: Remove the locks from SnapshotPurge and SnapshotSetProperty APIs.

This fixes the OM crash issue when the follower is lagging and it executes purgeKeys or snapshotMoveDeletedKeys for the Snapshot in one transaction.

Apache Jira: [HDDS-11137](#)

CDPD-71702: Ozone Manager is down to Snapshot Chain Corruption.

SSTFilteringService directly updates the snapshotInfoTable which can cause the snapshot chain corruption if OM crashes before DB gets flushed for snapshot purge and SSTFilteringService has updated the next snapshot in the chain.

Apache Jira: [HDDS-11068](#)

CDPD-71584: Ozone Recon DecommissioningInfo API throws displays the NPE error.

This issue is resolved by fixing the NullPointerException when running DecommissioningInfo API.

Apache Jira: [HDDS-11045](#)

CDPD-71502: Ozone Recon - Decommissioned datanodes show up even after removing it from the Recon Datanodes page.

Recon previously allowed to remove the Decommissioned datanodes and was removing from Recon rocksDB nodes table. However, Decommissioned datanodes continue to send heartbeats till they are being shutdown. This gets registered and added again in the Recon memory map which makes them show up again in datanodes UI. This issue is now resolved and allows only decommissioned datanodes to be removed and skip other node status or node operational status datanodes.

Apache Jira: [HDDS-11032](#)

CDPD-70469: Ozone Recon - Handle startup failure and log reasons as error because SCM non-HA is enabled.

This issue is now resolved by fixing the Recon startup failure when SCM runs in non-HA mode.

Apache Jira: [HDDS-10937](#)

CDPD-68912: Ozone Recon - Improve Recon startup failure handling.

This issue is now resolved. Recon should recover from Runtime or unexpected failures during startup and provide information on Recon UI. Recon can fail to start due to several reasons:

- Failure of registering of datanodes or invalid topology.
- Initialization of pipelines.

Apache Jira: [HDDS-10702](#)

CDPD-67668: Ozone Recon - Provide DN decommissioning detailed status and information inline with current CLI command output.

This issue resolved by adding a new improvement to provide API in Recon for DN decommissioning. Status and information is now inline with current CLI command output.

Apache Jira: [HDDS-10514](#)

CDPD-67460: Container Balancer should only move containers with size greater than 0 bytes.

This issue is now resolved by introducing a check on the size of the containers allowed to leave the source node during the balancing process.

Apache Jira: [HDDS-10483](#)

CDPD-67278: Fix the DN links on the Ozone SCM UI. This is a backport of KNOX-3012.

A change in Ozone affected Knox on the Ozone SCM UI. The links for the datanodes did not route through Knox. This issue is now resolved and the DN links will redirect to the correct Knox URLs.

CDPD-67095: DN URL in SCM Page through Knox redirects to non-Knox URL.

A change in Ozone affected Knox on the Ozone SCM UI. The links for the datanodes did not route through Knox. With CDPD-67278 and CDPD-69143, this issue is now resolved and the DN links will redirect to the correct Knox URLs.

CDPD-64874: Intermittent failure in TestOzoneRpcClientAbstract.testListSnapshot.

This issue is now resolved by fixing `listSnapshotAPI` intermittent wrong data issues. The `listSnapshot` API uses the `org.apache.hadoop.ozone.om.ListIterator.MinHeapIterator` which internally uses both `CacheIterator` and `DBIterator` and `DBIterator` had the logic of checking if rocks DB key is present in cache in `org.apache.hadoop.ozone.om.ListIterator.DbTableIter#getNextKey`. This checks the cache from table cache which may be intermittently flushed and makes the addition of duplicate entry in `org.apache.hadoop.ozone.om.ListIterator.MinHeapIterator`. You must use the pre-loaded keys in `org.apache.hadoop.ozone.om.ListIterator.CacheIter#cacheKeyMap` in `org.apache.hadoop.ozone.om.ListIterator.CacheIter`.

Apache Jira: [HDDS-9967](#)

CDPD-64815: NSSummary commands should close OzoneClient.

`NSSummaryAdmin` creates `OzoneClient` for some bucket-related checks. This issue now resolves:

- Close client when no longer needed
- Reuse client (or even bucket after lookup) for all checks

Apache Jira: [HDDS-9944](#)

CDPD-64209: Ozone Recon - Potential memory overflow in Container Health Task.

This issue is now resolved by fixing the Potential memory overflow in Container Health Task of Recon.

Apache Jira: [HDDS-9819](#)

CDPD-63596: Do not include SpotBugs at compile scope.

This issue is now resolved by removing spotbugs-annotation, an LGPL thirdparty dependency from the Ozone package.

Apache Jira: [HDDS-9692](#)

CDPD-62991: Recon UI - Bucket Drop down filter is not getting disabled when more than 1 volume is selected. This is a backport of HDDS-9556.

This issue is now resolved.

Apache Jira: [HDDS-9556](#)

CDPD-62931: Incorrect pipeline ID for closed container.

This issue is now resolved.

Apache Jira: [HDDS-9544](#)

CDPD-62925: Ozone debug chunkinfo command shows incorrect number of entries.

This issue is now resolved.

Apache Jira: [HDDS-9542](#)

CDPD-62471: Recon UI - Disk Usage page should reflect the information it displays.

This issue is now resolved.

Apache Jira: [HDDS-9465](#)

CDPD-62466: Improve thread names in Recon.

This issue is resolved by improving the thread naming in Recon process.

1. Pass Recon as a thread name prefix in Recon.
2. Ensure all other threads created in Recon code also include Recon in their name.

Apache Jira: [HDDS-9470](#)

CDPD-61700: Ozone debug chunkinfo shows incorrect block path for some nodes in a phatcat cluster.

This issue is now resolved.

Apache Jira: [HDDS-9356](#)

CDPD-60647: Snapshot purge should be an atomic operation.

This issue is resolved by fixing the OM crash issue when the follower is lagging and it executes purgeKeys or snapshotMoveDeletedKeys for the Snapshot in one transaction.

Apache Jira: [HDDS-9198](#)

CDPD-51724: SCM should avoid sending delete transactions for under-replicated containers.

This issue is now resolved.

Apache Jira: [HDDS-4368](#)

Fixed Issues in Apache Parquet

There are no fixed issues for Parquet in Cloudera Runtime 7.3.1.

Apache Patch Information

None

Fixed Issues in Phoenix

Review the list of Phoenix issues that are resolved in Cloudera Runtime 7.3.1.

OPSAPS-71553: OMID TSO server role fails during cluster installation

High Availability (HA) configuration generation is based on the number of nodes instead of a parameter in OMID.

OPSAPS-70507: ZDU runtime upgrade from 719 CHF4 to 7.3.0.1 fails for OMID

The OMID update issue is fixed.

OPSAPS-69838: OMID cannot connect contact the ZooKeeper cluster

OMID now connects to a secure port of ZooKeeper if TLS is enabled. Earlier it failed to communicate because port unification was missing.

OPSAPS-68583: ZooKeeper SSL/TLS support for OMID

OMID supports a secure connection to ZooKeeper if AutoTLS and ZooKeeper TLS are enabled.

CDPD-69649: The region server hosting the SYSTEM.CATALOG fails to serve any metadata requests as default handler pool threads are exhausted

This issue is fixed.

Apache Jira: [PHOENIX-6687](#)

CDPD-66730: Phoenix third party upgrade

Upgraded Guava to 32.0.1 due to CVE-2023-2976.

Fixed Issues in Ranger

Review the list of Ranger issues that are resolved in CDP Private Cloud Base 7.3.1.

CDPD-73663: RMS server threw ConcurrentModificationException

The original ConcurrentModificationException was likely thrown when the resource-mappings were modified in response to changes in the Hive metadata while they were being serialized for downloading to the NameNode (or secondary-namenode).

The fix is to create a shallow copy of resource-mappings before applying deltas which ensures that resource-mappings are not modified while they are being serialized for downloading to the NameNode.

CDPD-73326: Reduce memory needed to create Ranger policy engine

Ranger policy engine creates a RangerPolicyResourceMatcher object for every resource specified either in policy or in a tag association. PolicyResourceMatcher, for the services that have more than one level in their resource hierarchy, consists of RangerResourceMatcher objects for each level in the resource-level hierarchy for the resource. In many cases, this leads to creation of multiple RangerResourceMatchers with identical resource specification.

The fix for this issue avoids creation of multiple RangerResourceMatcher objects by maintaining a cache of them in the RangerPluginContext object associated with the Ranger policy engine, thereby reducing policy engine's memory needs.

CDPD-73144: Trie to support processing of evaluators during traversal

Ranger policy engine uses trie data structure to organize resources for faster retrieval of policies/tags/zones associated with a given resource. When a resource consists of multiple elements, like database/table/column, as many trie instances are consulted to retrieve policies/tags/zones associated with the resource. Such multi-trie retrieval can be optimized with a 2-pass traversal - first pass to get count and the second pass to get the actual objects. Trie data structure used in Ranger policy engine should be updated to support this optimization.

Now, Trie to support processing of evaluators during traversal is enhanced.

CDPD-72207: ll_service_id is empty for an invalid notification type

The fix corrects the query to fetch latestInvalidNotificationId even though ll_service_id is empty. This ensures that NameNode gets the appropriate delta's mappings.

CDPD-72203: Users observing role change from ROLE_SYS_ADMIN to ROLE_USER

Fixes role reset (to USER role) for users in usersync paged requests to ranger-admin.

CDPD-71719: Ranger override policy was not working

Ranger override policy was not allowing the access even though all permissions were given to the user.

This fix ensures that once all of the requested accesses are successfully allowed by (possibly multiple) Ranger policies, the access evaluation terminates with access allowed as the result.

CDPD-70081: "Drop database cascade" resulted in dropping of a table on which the user did not have access

Drop database cascade failed if the user did not have access to one or more of the underlying tables. It deleted the tables the user had access to but not others which caused the database to be not dropped as well.

This issue is fixed now.

CDPD-70003: Ranger KMS solr auditing fails when secure zk port 2182 is used

The fix includes the netty specific libs so that Ranger KMS to Solr supports ZooKeeper-SSL enabled connection.

CDPD-69488: Upgrade failure due to NPE in PatchForUpdatingServiceDefJson_J10058

Patch upgrade error failure in non-default service-def is fixed now.

CDPD-69305: /plugins/policies/importPoliciesFromFile API returns 500 service connectivity error through Knox Proxy

The fix imports large policy files using the Ranger importPoliciesFromFile API through Knox.

CDPD-68921: Exclude flag not taking effect for Ozone key resource in Ranger policy

Fix for exclude flag not taking effect for Ozone key resource in Ranger policy has been added.

CDPD-67823: Ranger RMS gives all permissions to the user through the Create permission

An additional check is now made to ensure that the user attempting to alter a HDFS directory that maps to the Hive database is owner of the Hive database for the attempted operation is allowed.

CDPD-67193: Issue with inactivityTimeout getting reset

The inactivityTimeout was getting reset when a user updated its profile from the UserProfile page.

Fixed issue of not resetting inactivityTimeout to a default value of 15 minutes when user updates its profile from UserProfile page on Ranger Admin UI.

CDPD-66927: HDFS authorization logic for directory hierarchy rooted at "/" is incorrect

Ranger authorization logic for the HDFS commands that require authorization of the entire directory hierarchy rooted at the specified directory argument is incorrect as it does not correctly compute the sub-directory paths. The paths of sub-directories that need to be authorized incorrectly contain an extra '/' character, which leads to incorrect authorization results.

The issue is fixed now.

CDPD-66842: Ranger Admin server gives empty response

Ranger Admin server gave an empty response when a user with user-role tried to update lastname or email address.

The issue is fixed now. Error response with message will be shown when a user with user-role tries to add/update last name or email address.

CDPD-66839: Enhance perf-tracer to get CPU time when possible

Ranger module is instrumented with performance measurement code. It enables performance logging for the module and helps in measuring the amount of time spent during execution of various methods/functions during its operation. For achieving more precise time measurement, this feature supports nanosecond precision when the JVM version supports it.

CDPD-66624: Transform URLs with or without “/” at the end issue

The fix enables the transformation step handle “/” at the end of the path.

CDPD-66404: Merging apache ranger jiras for handling local storage data for column show/hide functionality

Implemented Column Hide/Show functionality in Audit Plugin Status tab.

CDPD-66358: HS2 logs having a huge number of WARN logs

HS2 logs had a huge number of WARN logs from RangerHiveAuthorizer regarding connection to HMS for fetching Hive object owner.

This fix addresses the issue where HS2 logs have a huge number of WARN logs.

CDPD-66136: Display of query information for Show databases/schemas command on Ranger Admin UI

In Ranger React UI, if the resource type for certain commands were logged as "null" in the audits, then in the access audits, the information of the query/operations performed would not be displayed.

This ticket addresses the issue and displays the query/operation information for access audits where the resource type was "null".

CDPD-66092: Ranger Javapatch failure even if service-defs do not exist in Ranger DB

Added support to upgrade non-default service-defs in Ranger.

CDPD-65923: Audit logs for Mask and Row policy does not show policy condition under policy item

The fix now shows policy conditions under policy items for Mask and Row policy Audit logs.

CDPD-65650: Pagination missing on the Ranger Admin - Plugin Status page

This fix offers the following:

- Sorting works properly after this patch.
- Pagination added.

CDPD-63747: Cache the results of access evaluation

This feature trades off more memory requirement against a potential faster evaluation of policies when chained-plugin (as when RMS is enabled) is configured for HDFS storage authorization. If the configuration parameter "ranger.plugin.hdfs.useResultCache" (default:false) is set to true, then the result of Hive policy authorization for a HDFS storage location is cached and is reused in subsequent accesses of that HDFS location.

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

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

OPSAPS-69411: Update AuthzMigrator GBN to point to latest non-expired GBN

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

Apache Patch information

- RANGER-4973
- RANGER-4972
- RANGER-4960
- RANGER-4933

- RANGER-4912
- RANGER-4905
- RANGER-4893
- RANGER-4833
- RANGER-4823
- RANGER-4819
- RANGER-4818
- RANGER-4802
- RANGER-4799
- RANGER-4798
- RANGER-4797
- RANGER-4796
- RANGER-4791
- RANGER-4786
- RANGER-4782
- RANGER-4781
- RANGER-4780
- RANGER-4774
- RANGER-4767
- RANGER-4753
- RANGER-4747
- RANGER-4745
- RANGER-4737
- RANGER-4729
- RANGER-4722
- RANGER-4720
- RANGER-4718
- RANGER-4717
- RANGER-4710
- RANGER-4699
- RANGER-4698
- RANGER-4690
- RANGER-4689
- RANGER-4688
- RANGER-4681
- RANGER-4673
- RANGER-4668
- RANGER-4653
- RANGER-4641
- RANGER-4611
- RANGER-4609
- RANGER-4607
- RANGER-4598
- RANGER-4597
- RANGER-4596
- RANGER-4595
- RANGER-4594
- RANGER-4593
- RANGER-4591
- RANGER-4590

- RANGER-4589
- RANGER-4588
- RANGER-4586
- RANGER-4578
- RANGER-4577
- RANGER-4576
- RANGER-4575
- RANGER-4574
- RANGER-4573
- RANGER-4568
- RANGER-4555
- RANGER-4554
- RANGER-4553
- RANGER-4552
- RANGER-4551
- RANGER-4550
- RANGER-4549
- RANGER-4548
- RANGER-4547
- RANGER-4546
- RANGER-4545
- RANGER-4544
- RANGER-4532
- RANGER-4515
- RANGER-4513
- RANGER-4492
- RANGER-4370
- RANGER-4303
- RANGER-4278
- RANGER-4261
- RANGER-4229
- RANGER-4221
- RANGER-4172
- RANGER-4010
- RANGER-3805
- RANGER-3772
- RANGER-3759
- RANGER-3745
- RANGER-3657
- RANGER-3182
- RANGER-3174

Fixed Issues in Ranger KMS

Review the list of Ranger KMS issues that are resolved in CDP Private Cloud Base 7.3.1.

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

KEYTRUSTEE_SERVER and RANGER_KMS_KTS services are not supported starting from the CDP 7.3.1.0 release. Therefore added validation and confirmation messages to the CM upgrade

wizard to alert the user to migrate KEYTRUSTEE_SERVER keys to RANGER_KMS before upgrading to CDP 7.3.1.0 release.

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

The Keytrustee components - KEYTRUSTEE_SERVER and RANGER_KMS_KTS services are not supported starting from the CDP 7.3.1.0 release. These services cannot be installed or managed with CM 7.13.1.0 using CDP 7.3.1.0.

CDPD-19186: Replacement of algorithm PBKDF2WithMD5AndTripleDES for Ranger KMS key operations

Support for PBKDF2WithHmacSHA256 is added in KMS.

Code to decrypt the Masterkey and all Zonekeys using the older algorithm and then re-encrypt it using the latest algorithm is implemented.

Fixed Issues in Schema Registry

Review the list of Schema Registry issues that are resolved in Cloudera Runtime 7.3.1.

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

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

CDPD-60160: Schema Registry Atlas integration does not work with Oracle databases

The Schema Registry Atlas integration works correctly when Oracle is used as the database of Schema Registry.

CDPD-58990: The getSortedSchemaVersions method should order by version number and not by schemaVersionId

Schemas are ordered correctly by their version number during validation instead of their ID number.

CDPD-58949: Schemas are deduplicated during import

Importing works correctly and Schema Registry does not deduplicate imported schema versions.

CDPD-59015: Schema Registry does not create new versions of schemas even if the schema is changed

The new fingerprinting version solves this issue.

Fixed Issues in Apache Solr

Review the list of Solr issues that are resolved in Cloudera Runtime 7.3.1.

OPSAPS-71690: Update control group V2 configuration parameters

The default values of the control group (CGroup) V2 configuration parameters are updated in Cloudera Manager for the Solr service. The following table describes the default values of the corresponding V2 parameters.

Parameter name	Default values
memory.high	-1
memory.max	-1
io.weight	100
cpu.weight	100

For more information on CGroup V2 parameters, see [Configuring Resource Parameters](#).

Fixed Issues in Spark

Review the list of Spark issues that are resolved in Cloudera Runtime 7.3.1.

CDPD-74697 - Spark Iceberg vectorized Parquet read of decimal column is incorrect

CDPD-72774 - Use common versions of commons-dbcp2 and commons-pool2

CDPD-70114 - Redirect spark-submit, spark-shell etc. scripts to their Spark 3 counterparts

CDPD-58844 - Spark - Upgrade Janino to 3.1.10 due to CVE-2023-33546

CDPD-48171 - Spark3 - Upgrade snakeyaml due to CVE-2022-1471

Apache patch information

None

Fixed Issues in Streams Messaging Manager

Review the list of Streams Messaging Manager (SMM) issues that are resolved in Cloudera Runtime 7.3.1.

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

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

CDPD-72543: Security headers are not set for static files in SMM

SMM now applies the following security-related headers to static files:

- Content-Security-Policy
- X-XSS-PROTECTION
- X-Content-Type-Options
- X-Frame-Options
- Strict-Transport-Security

CDPD-73643: Unused CM_USER parameter is visible in /cm-configs internal endpoint

The unused CM_USER field has been removed from the /cm-configs internal endpoint

CDPD-70313: KNOX does not send Authentication header on FIPS configuration

KNOX now sends the Authentication header on FIPS clusters.

CDPD-53437: Alert notification messages do not contain an alert cause

In some cases, the resources and the attributes of the resource that triggered the alert were not included in the alert message. Typically, this happened when the alert policy defined a WITH ANY clause. From now on, resources that trigger the alert are always included in the alert message.

CDPD-61516: Anchor links are broken on Knox enabled clusters

Links that appear on the Overview page when hovering over producers, consumers, or partitions are no longer broken on Knox-enabled clusters.

Fixed Issues in Streams Replication Manager

Review the list of Streams Replication Manager (SRM) issues that are resolved in Cloudera Runtime 7.3.1.

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

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

Fixed Issues in Apache Tez

There are no fixed issues for Tez in Cloudera Runtime 7.3.1.

Fixed Issues in YARN and YARN Queue Manager

Review the list of YARN and YARN Queue Manager issues that are resolved in Cloudera Runtime 7.3.1.

COMPX-17702: Backport - YARN-10345 - HsWebServices containerlogs does not honor ACLs for completed jobs

The following rest APIs now have ACL authorization:

- /ws/v1/history/containerlogs/{containerid}/{filename}
- /ws/v1/history/containers/{containerid}/logs

Apache Jira: [YARN-10345](#)

COMPX-16285: Optimize system credentials sent in node heartbeat responses

Previously, the heartbeat responses set all application's tokens even though all applications were not active on a node. Hence, for each node and each heartbeat too many SystemCredentialsFor AppsProto objects were created. This issue is now resolved and the system credentials sent in node heartbeat responses are optimized..

Apache Jira: [YARN-6523](#)

CDPD-73754: Yarn Application Master Node web link is broken on yarnui2 page

Previously, the RM did not open the Yarn application manager node web link on the **yarnui2** page because the URL ended with a /. This issue is now resolved and the last character / is now removed from the URL.

Apache Jira: [YARN-11729](#)

Fixed Issues in Zookeeper

Review the list of ZooKeeper issues that are resolved in Cloudera Runtime 7.3.1.

CDPD-67821: Zookeeper - Information disclosure in persistent watcher handling(CVE-2024-23944)

There was information disclosure in persistent watchers handling in Apache ZooKeeper due to CVE-2024-23944. This issue is now fixed.

Apache Jira: [ZOOKEEPER-4799](#)

CDPD-66977: Backport ZOOKEEPER-4804 Use daemon threads for Netty client

Previously, when the Netty client was used, the Java process did not respond on System.exit when the Zookeeper connection was open. This issue was caused by the non-daemon threads created by Netty. This issue is now resolved.

Apache Jira: [ZOOKEEPER-4804](#)

Known Issues In Cloudera Runtime 7.3.1

This topic describes known issues and workarounds in this release of Cloudera Runtime.

Known Issues in Apache Atlas

Learn about the known issues in Atlas, the impact or changes to the functionality, and the workaround.

CDPD-67112: Import transforms do not work as expected when replacing a string which already has ":"

The character ":" is not supported in path replacements. The import succeeds but location remains unchanged. The character ":" must be avoided.

None

CDPD-69150: Unable to add labels or user defined properties in Japanese

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

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-68191: Suggestions do not return the correct results when searching multiple Chinese characters

Free text search does not return results when searching Chinese phrases made of multiple characters. Partial searches return the correct results.

None.

CDPD-71219: Regression : Suggestions don't work for single character words on indexable attributes

When searching for entities whose name (entity names are indexable) is a single character, search results are returned but suggestions are not. When searching for entities whose description (entity descriptions are not indexable) is a single character, both search results and suggestions are returned.

None.

CDPD-67450: Table name renaming operation is not updating or creating iceberg_table entity

Renaming an Iceberg table does not update the corresponding Atlas entity.

None.

CDPD-67089: Export/Import: When a table with Ozone path is exported as "connected", only the Ozone key is exported.

When table with Ozone path is exported as "connected", only Ozone key is exported. Other Ozone entities, such as Ozone volume, Ozone bucket are not exported.

None.

CDPD-59565: Whole lineage becomes hidden when filters are enabled with deleted entity

When a larger lineage structure contains a circular lineage and this circular lineage contains at least one deleted entity, the whole lineage structure becomes hidden if both of the following filters are used:

- Hide Processes
- Hide Deleted Entities

None.

CDPD-43772: Performance issues with Atlas service

If there are lot of update operations and the compression type of column families of `atlas_janus` table is SNAPPY, then the Kafka message processing might become slower.

- Consider setting compression type of column families of `atlas_janus` table as GZ.

CDPD-69910: NPE while deleting BusinessMetadata

If business metadata is created without adding any applicable types, a `NullPointerException` is produced when we try to delete that business metadata.

None

Apache Jira: [ATLAS-4863](#)

CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeupException: null" is seen

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

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-74180: Export/Import : If Shell entities have a lineage, it is not exported

If there is a shell entity which has lineage, while using the Export API, that shell entity will not be exported in the zip file.

CDPD-66938: [Analyze] [Atlas] [FIPS] test_time_range tests fail

When the Apache Atlas server is running on a node which has time zone other than UTC, there might be a time of day when the search results might differ if the relative **CreateTime** date range filters of TODAY, YESTERDAY, etc. are used.

Use explicit date range filters instead of using relative date range filters, such as, TODAY, YESTERDAY.

CDPD-70321: Atlas Parallel import is failing with various errors

During a parallel import-export activity with six iceberg table policies with `exportOption` as `db1.*` for all six exports, all import fail after the exports.

CDPD-76269: POST Rolling Upgrade performed from 7.1.7.3000 to 7.3.1.0 , and then downgraded from 7.3.1.0 to 7.1.7.3000 , updating edge to enable tag propagation is failing

When performing a rolling upgrade from 7.1.7.3000 to 7.3.1.0, and then downgrading from 7.3.1.0 to 7.1.7.3000, updating edge to enable tag propagation fails.

Tag propagation works well after a restart.

Known Issues in Apache Avro

Learn about the known issues in Avro, the impact or changes to the functionality, and the workaround.

CDPD-23451: Avro library depends on the already EOL jackson-mapper-asl 1.9.13-cloudera.1 that also contains a couple of CVEs. The jackson library is part of the Avro API so cannot be changed without a complete rebase.

None.

Known Issues in Cloud Connectors

Learn about the known issues in Cloud Connectors, the impact or changes to the functionality, and the workaround.

AWS SDK 2.25.53 warning about transfer manager not using CRT client

Due to the AWS SDK 2.25.53 upgrade, the following warning might be seen:

```
5645:2024-09-13 16:29:17,375 [setup] WARN s3.S3TransferManager
```

```
(LoggerAdapter.java:warn(225)) - The provided S3AsyncClient
is an instance of
    MultipartS3AsyncClient, and thus multipart download fe
ature is not enabled. To benefit
    from all features, consider using S3AsyncClient.crtBu
ilder().build() instead
```

This error message is completely harmless and should be ignored. For more information, see [HADOOP-19272](#).

None

Known issues in Cruise Control

Learn about the known issues in Cruise Control, the impact or changes to the functionality, and the workaround.

CDPD-44676: Rebalancing with Cruise Control does not work if the metric reporter fails to report the CPU usage metric

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.CruiseCont
rolMetricsReporter:
    [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 Hadoop

There are no known issues for Hadoop in Cloudera Runtime 7.3.1.

Known Issues in Apache HBase

Learn about the known issues in HBase, the impact or changes to the functionality, and the workaround.

CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala

- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

OpDB Data Hub cluster fails to initialize if you are reusing a cloud storage location that was used by an older OpDB Data Hub cluster

Workaround: Stop HBase using Cloudera Manager before deleting an OpDB Data Hub cluster.

IntegrationTestReplication fails if replication does not finish before the verify phase begins

During IntegrationTestReplication, if the verify phase starts before the replication phase finishes, the test will fail because the target cluster does not contain all of the data. If the HBase services in the target cluster does not have enough memory, long garbage-collection pauses might occur.

Workaround: Use the -t flag to set the timeout value before starting verification.

HDFS encryption with HBase

Cloudera has tested the performance impact of using HDFS encryption with HBase. The overall overhead of HDFS encryption on HBase performance is in the range of 3 to 4% for both read and update workloads. Scan performance has not been thoroughly tested.

Workaround: N/A

Snappy compression with /tmp directory mounted with noexec option

Using the HBase client applications such as hbase hfile on the cluster with Snappy compression could result in UnsatisfiedLinkError.

Add -Dorg.xerial.snappy.tmpdir=/var/hbase/snappy-tmpdir to Client Java Configuration Options in Cloudera Manager that points to a directory where exec option is allowed.

AccessController postOperation problems in asynchronous operations

When security and Access Control are enabled, the following problems occur:

- If a Delete Table fails for a reason other than missing permissions, the access rights are removed but the table may still exist and may be used again.
- If hbaseAdmin.modifyTable() is used to delete column families, the rights are not removed from the Access Control List (ACL) table. The portOperation is implemented only for postDeleteColumn().
- If Create Table fails, full rights for that table persist for the user who attempted to create it. If another user later succeeds in creating the table, the user who made the failed attempt still has the full rights.

Workaround: N/A

Apache Issue: [HBASE-6992](#)

HBase shutdown can lead to inconsistencies in META

Cloudera Manager uses an incorrect shutdown command. This prevents graceful shutdown of the HBase service and forces Cloudera Manager to kill the processes instead. It can lead to inconsistencies in Meta.

Workaround: Run the following command instead of shutting down the HBase service using Cloudera Manager.

```
hbase master stop --shutDownCluster
```

The command output must end with Closing master protocol: MasterService phrase. You can verify the command execution by checking the master logs. The log must contain Cluster shutdown requested of master=xxx and the closing of regions. Upon successful execution, the RegionServers start shutting down.



Note: The command does not stop the *REST Server* and the *Thrift Server* role instances. You can safely shut down them from Cloudera Manager later.

If you find any inconsistencies, please contact Cloudera Support.

Bulk load is not supported when the source is the local HDFS

The bulk load feature (the `completebulkload` command) is not supported when the source is the local HDFS and the target is an object store, such as S3/ABFS.

Workaround: Use `distcp` to move the HFiles from HDFS to S3 and then run bulk load from S3 to S3.

Apache Issue: N/A

Storing Medium Objects (MOBs) in HBase is currently not supported

Storing MOBs in HBase relies on bulk loading files, and this is not currently supported when HBase is configured to use cloud storage (S3).

Workaround: N/A

Apache Issue: N/A

Known Issues in HDFS

Learn about the known issues in HDFS, the impact or changes to the functionality, and the workaround.

CDPD-65530: HDFS requests throw `UnknownHostException` during OS upgrade

During the VM replacement as part of OS upgrade, every new node gets a new IP Address, and if the old IP address is cached somewhere, HDFS requests fail with `UnknownHostException` and it recovers after sometime (10 mins max).

The issue is seen during COD and DL ZDU.

None.

CDPSDX-5302: Avoiding long delay on the HBase master does not happen during upgrade.

1. Log in to Cloudera Manager
2. Select the HDFS service
3. Select Configurations tab
4. Search for `hdfs-site.xml`.
5. Set `ipc.client.connect.timeout = 5000`
6. Set `ipc.client.connect.max.retries.on.timeouts = 5`
7. Click Save

The above configuration changes ensures that:

1. The long delay on the HBase master does not happen during upgrade.
2. The long delay on the HBase master recovery does not happen during upgrade.

CDPD-67230: Rolling restart can cause failed writes on small clusters

In a rolling restart, if the cluster has less than 10 datanodes existing writers can fail with an error indicating a new block cannot be allocated and all nodes are excluded. This is because you have attempted to use all the datanodes in the cluster, and failed to write to each of them as they were restarted. This only happen on small clusters of less than 10 datanodes, because larger clusters have more spare nodes to allow the write to continue.

None.

CDPD-60873: java.io.IOException:Encountered "status=ERROR, status message, ack with firstBadLink" while fixing the HDFS corrupt file during rollback.

Increase the value of dfs.client.block.write.retries to the number of nodes in the cluster and perform Deploy client configuration procedure for rectification.

CDPD-60431: Configuration difference between 7.1.7 SP2 and 7.1.9.0 results

Component	Configuration	Old Value	New Value	Description
HDFS	dfs.permissions.ContentSummary.subAccess	Not Set	True	Performance optimization for NameNode content summary API
HDFS	dfs.datanode.handler.count	8	10	Optimal value for DN server threads on large clusters

None.

CDPD-60387: Configuration difference between 7.1.8.3 and 7.1.9.0 results

Component	Configuration	Old Value	New Value	Description
HDFS	dfs.namenode.access.precision	None	0	Optimal value for NameNode performance on large clusters
HDFS	dfs.datanode.handler.count	8	10	Optimal value for DN server threads on large clusters

None.

OPSAPS-64307: When the JournalNodes on a cluster are restarted, the Add new NameNode wizard for HDFS service might fail to bootstrap the new NameNode. If there was no new fsImage created from the time JournalNodes restarted, during the restart the edit logs were rolled in the system.

If the bootstrapping fails during the Add new NameNode wizard, then perform the following steps:

1. Delete the newly added NameNode and FailoverController
2. Move the active HDFS NameNode to safe mode
3. Perform the Save Namespace operation on the active HDFS NameNode
4. Leave safe mode on the active HDFS NameNode
5. Add the new NameNode again



Note: Entering safe mode disables writes to HDFS which causes a service disruption. If you cannot enter the safe mode, delete the newly added NameNode and FailoverController in the HDFS service and wait until HDFS automatically creates a new fsImage and then add the new NameNode again with the wizard.

OPSAPS-64363: Deleting of additional Standby Namenode does not delete the ZKFC role and this has to be done manually.

None.

CDPD-28390: Rolling restart of the HDFS JournalNodes may time out on Ubuntu20.

If the restart operation times out, you can manually stop and restart the Name Node and Journal Node services one by one.

OPSAPS-55788: WebHDFS is always enabled. The Enable WebHDFS option does not take effect.

None.

OPSAPS-63299: Disable HA command for a nameservice does not work if the nameservice has more than 2 NameNodes defined.

None.

OPSAPS-63301: Deleting nameservice command does not delete all the NameNodes belonging to the nameservice, if there are more than two NameNodes that are assigned to the nameservice.

None.

Unsupported features

The following HDFS features are currently not supported in Cloudera Data Platform:

- ACLs for the NFS gateway ([HADOOP-11004](#))
- Aliyun Cloud Connector ([HADOOP-12756](#))
- Allow HDFS block replicas to be provided by an external storage system ([HDFS-9806](#))
- Consistent standby Serving reads ([HDFS-12943](#))
- Cost-based RPC FairCallQueue ([HDFS-14403](#))
- HDFS Router Based Federation ([HDFS-10467](#))
- NameNode Federation ([HDFS-1052](#))
- NameNode Port-based Selective Encryption ([HDFS-13541](#))
- Non-Volatile Storage Class Memory (SCM) in HDFS Cache Directives ([HDFS-13762](#))
- OpenStack Swift ([HADOOP-8545](#))
- SFTP FileSystem ([HADOOP-5732](#))
- Storage policy satisfier ([HDFS-10285](#))

Technical Service Bulletins

TSB 2022-549: Possible HDFS Erasure Coded (EC) data loss when EC blocks are over-replicated

Cloudera has detected a bug that can cause loss of data that is stored in HDFS Erasure Coded (EC) files in an unlikely scenario.

Some EC blocks may be inadvertently deleted due to a bug in how the NameNode chooses excess or over-replicated block replicas for deletion. One possible cause of over-replication is running the HDFS balancer soon after a NameNode goes into failover mode.

In a rare situation, the redundant blocks can be placed in such a way that one replica is in one rack, and few redundant replicas are in the same rack. Such placement causes a counting bug ([HDFS-16420](#)) to be triggered. Instead of deleting just the redundant replicas, the original replica may also be deleted.

Usually this is not an issue, because the lost replica can be detected and reconstructed from the remaining data and parity blocks. However, if multiple blocks in an EC Block Group are affected by this counting bug within a short time, the block cannot be reconstructed anymore. For example, 4 blocks are affected out of 9 for the RS(6,3) policy.

Another situation is recommissioning multiple nodes back into the same rack of the cluster where the current live replica exists.

Upstream JIRA

[HDFS-16420](#)**Knowledge article**

For the latest update on this issue see the corresponding Knowledge article: [TSB 2022-549: Possible HDFS Erasure Coded \(EC\) data loss when EC blocks are over-replicated](#)

Known Issues in Apache Hive

Learn about the known issues in Hive, the impact or changes to the functionality, and the workaround.

DAG not retried after failure

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 Hue

Learn about the known issues in Hue, the impact or changes to the functionality, and the workaround.

OPSAPS-69659: Hue service fails on restart with "Unable to find pycopg2 2.5.4" error

Hue service fails to restart and you see the following error: Unable to find pycopg2 2.5.4. This could be because you have installed Python in a non-default location and Hue is unable to locate the pycopg2 PostgreSQL database adapter.

You must specify the path where you have installed Python in the PYTHONPATH property in the Hue Advanced Configuration Snippet using Cloudera Manager.

1. Log in to Cloudera Manager as an Administrator.
2. Go to Clusters Hue Configurations and add the following key and value in the Hue Service Environment Advanced Configuration Snippet (Safety Valve) field:

Key: PYTHONPATH

Value: `[***PYTHON-PATH***]`

Replace `[***PYTHON-PATH***]` with the actual location where you have installed Python. For example, `/opt/cloudera/parcels/CDH/lib/hue/build/env/lib/python3.8/site-packages`

3. Click Save Changes.
4. Restart the Hue service.

CDPD-58978: Batch query execution using Hue fails with Kerberos error

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 Hue and this is not supported on a secure cluster.

There is no workaround. You can submit the queries individually.

CDPD-59677: Unable to view Phoenix tables on the left assist in Hue

On clusters secured with Knox, you may not be able to see Phoenix tables on the left assist that are present under the default database (that is, an empty("") database).

None.

CDPD-58142: A query is not pre-populated in the Hue editor after clicking on the "Re Execute" button

When you click Re Execute to rerun a query from the Job Browser Queries Query Details page, the query does not get populated on the Hue editor, as expected.

None.

CDPD-54376: Clicking the home button on the File Browser page redirects to HDFS user directory

When you are previewing a file on any supported filesystem, such as S3 or ABFS, and you click on the Home button, you are redirected to the HDFS user home directory instead of the user home directory on the said filesystem.

None.

CDPD-41306: pip3.8 freeze command does not work and results into an error

You see the `/usr/lib/hue/build/env/bin/python: No such file or directory` error when you run the following command:

```
build/env/bin/pip3.8 freeze
```

Run the freeze command as follows by specifying the paths of python3.8 and pip3.8:

```
/opt/cloudera/parcels/CDH/lib/hue/build/env/bin/python3.8 /opt/cloudera/parcels/CDH/lib/hue/build/env/bin/pip3.8 freeze
```

CDPD-43293: Unable to import Impala table using Importer

Creating Impala tables using the Hue Importer may fail.

If you have both Hive and Impala services installed on your cluster, then you can import the table using by selecting the Hive dialect from **Tables Sources**.

If only Impala service is installed on your cluster, then go to **Cloudera Manager Clusters Hue Configurations** and add the following line in the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` field:

```
[beeswax]
max_number_of_sessions=1
```

CDPD-41136: Importing files from the local workstation is disabled by default

Cloudera has disabled the functionality to import files from your local workstation into Hue because it may cause errors. You may not see the Local File option in the Type drop-down menu on the **Importer** page by default.

You can enable the functionality to import files from your local workstation by specifying the following parameter in the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` field in Hue configurations in Cloudera Manager:

```
[indexer]
enable_direct_upload=true
```

DWX-8602: Unable to import a large CSV file from the local workstation

You may see an error message while importing a CSV file into Hue from your workstation, stating that you cannot import files of size more than 200 KB.

Upload the file to S3 or ABFS and then import them into Hue using the Importer.

CDPD-39330: Unable to use the pip command in CDP

You may not be able to use the pip command in CDP 7.1.7 or higher and may see the following error when using pip in a command: "ImportError: cannot import name chardet".

Follow the steps listed on [Unable to use pip command in CDP](#).

CDPD-24294: Hue uses the unsafe-inline directive in its Content Security Policy (CSP) header

Hue 4 web interface uses the unsafe-inline directive in its CSP header. As a result, the application server does not set the CSP header in its HTTP responses, and therefore does not benefit from the additional protection against potential cross-site scripting issues and other modern application

vulnerabilities which a properly configured CSP may provide. This could lead to application vulnerability.

Cloudera recommends deploying additional security measures such as a firewall within the Hue server to control allowed connections, and SSO-based authentications mechanisms such as LDAP or SAML.

OPSAPS-61244: Cloudera Manager displays stale Hue configuration after upgrading to CDP 7.1.x from CDH 6.

After upgrading from CDH 6 to CDP 7.1.x, you may see stale configurations in Cloudera manager for the Hue service.

Manually restart the Hue service from Cloudera Manager.

ENGESC-9091: Setting idle session timeout for Hue does not work when the cluster is secured using Knox SSO

If Hue is configured with `desktop.auth.backend.KnoxSpnegoDjangoBackend` as the Authentication Backend, then the automatic idle session logout that is set by configuring the `idle_session_timeout` property does not take effect. You may also see 404 error while accessing Hue from the Knox UI when the `idle_session_timeout` property is not set to -1.

None

DOCS-10377: Hue UI is blank upon login after upgrading to CDP 7.1.7 from CDH 6

If your cluster was secured using Knox, and if you have upgraded from CDH 6 to CDP 7.1.7, then you may see a blank Hue screen. This could happen because the `knox_proxyhosts` parameter is newly introduced in CDP, and it is possible that this parameter is not configured in Cloudera Manager under Hue configuration.

Specify the host on which you have installed Knox in the Hue Knox Proxy Hosts configuration as follows:

1. Log in to Cloudera Manager as an Administrator.
2. Obtain the host name of the Knox Gateway by going to Clusters Knox service Instances .
3. Go to Clusters Hue service Configuration and search for the Knox Proxy Hosts field.



Note: Cloudera Manager displays the following warning if the Knox Proxy Hosts field is empty when Knox Gateway is enabled on the CDP cluster: The parameter `knox_proxyhosts` cannot be empty. This can happen if there are no Knox Gateways. Please set the `knox_proxyhosts` to the list of hosts that have Knox Gateways.

4. Specify the Knox Gateway hostname in the Knox Proxy Hosts field.
5. Click Save Changes and restart the Hue service.

OPSAPS-58927: Connection failed error when accessing the Search app (Solr) from Hue

If you are using Solr with Hue to generate interactive dashboards and for indexing data, and if you have deployed two Solr services on your cluster and selected the second one as a dependency for Hue, then Cloudera Manager assigns the hostname of the first Solr service and the port number of the second Solr service generating an incorrect Solr URL in the search section of the `hue.ini` file. As a result, you may see a “Connection failed” error when you try to access the Search app from the Hue web UI.

1. Log into Cloudera Manager as an Administrator.
2. Go to Clusters Hue service Configuration and add the following lines in the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` field:

```
[search]
# URL of the Solr Server
```

```
solr_url=http://[***HOSTNAME***]:[***PORT***]/solr/
```

For example:

```
solr_url=http://solr2:4567/solr/
```

3. Click Save Changes.
4. Restart the Hue service.

CLR-72255: Error while rerunning Oozie workflow

You may see an error such as the following while rerunning an already executed and finished Oozie workflow through the Hue web interface: E0504: App directory [hdfs://cdh/user/hue/oozie/workspaces/hue-oozie-1571929263.84] does not exist.

To resolve this issue, add the following property in the Hue Load Balancer Advanced Configuration Snippet:

1. Sign in to Cloudera Manager as an administrator.
2. Go to Clusters Hue service Configurations Load Balancer and search for the Load Balancer Advanced Configuration Snippet (Safety Valve) for httpd.conf field.
3. Specify MergeSlashes OFF in the Load Balancer Advanced Configuration Snippet (Safety Valve) for httpd.conf field.
4. Click Save Changes.
5. Restart the Hue Load Balancer.

CDPD-16407: Python-psycopg2 package version 2.8.4 not compatible with Hue

Ubuntu 18.04 provides python-psycopg2 package version 2.8.4 but it is not compatible with Hue because of a bug in the Django framework.

Downgrade the package at the OS level by running the following command:

```
sudo apt install python-psycopg2==2.7.5
```

or install python-psycopg2 package using pip by running the following command:

```
sudo pip install psycopg2==2.7.5
```

DOCS-6344: Hue limitation after upgrading from CDH to CDP Private Cloud Base

The hive.server2.parallel.ops.in.session configuration property changes from TRUE to FALSE after upgrading from CDH to CDP Private Cloud Base. Current versions of Hue are compatible with this property change; however, if you still would like to use an earlier version of Hue that was not compatible with this property being FALSE and shared a single JDBC connection to issue queries concurrently, the connection will no longer work after upgrading.

CDPD-43956: Manually replace UUID when importing Oozie workflows containing sub-workflows

When importing Oozie workflows that contain sub-workflows, you must replace all the UUID entries with unique new entries. If you change the UUID of a sub-workflow, you must update that reference in the parent workflow to avoid circular dependencies.

INSIGHT-3707: Query history displays "Result Expired" message

You see the "Result Expired" message under the Query History column on the **Queries** tab for queries which were run back to back. This is a known behaviour.

None.

Unsupported features

CDPD-59595: Spark SQL does not work with all Livy servers that are configured for High Availability

SparkSQL support in Hue with Livy servers in HA mode is not supported. Hue does not automatically connect to one of the Livy servers. You must specify the Livy server in the Hue Advanced Configuration Snippet as follows:

```
[desktop]
[spark]
livy_server_url=http(s)://[***LIVY-FOR-SPARK3-SERVER-HOST***]:
[***LIVY-FOR-SPARK3-SERVER-PORT***]
```

Moreover, you may see the following error in Hue when you submit a SparkSQL query: Expecting value: line 2 column 1 (char 1). This happens when the Livy server does not respond to the request from Hue.

Specify all different Livy servers in the `livy_server_url` property one at a time and use the one which does not cause the issue.

CDPD-18491: PySpark and SparkSQL are not supported with Livy in Hue

Hue does not support configuring and using PySpark and SparkSQL with Livy in CDP Private Cloud Base.

Importing and exporting Oozie workflows across clusters and between different CDH versions is not supported

You can export Oozie workflows, schedules, and bundles from Hue and import them only within the same cluster if the cluster is unchanged. You can migrate bundle and coordinator jobs with their workflows only if their arguments have not changed between the old and the new cluster. For example, hostnames, NameNode, Resource Manager names, YARN queue names, and all the other parameters defined in the `workflow.xml` and `job.properties` files.

Using the import-export feature to migrate data between clusters is not recommended. To migrate data between different versions of CDH, for example, from CDH 5 to CDP 7, you must take the dump of the Hue database on the old cluster, restore it on the new cluster, and set up the database in the new environment. Also, the authentication method on the old and the new cluster should be the same because the Oozie workflows are tied to a user ID, and the exact user ID needs to be present in the new environment so that when a user logs into Hue, they can access their respective workflows.



Note: Migrating Oozie workflows from HDP clusters is not supported.

Known Issues in Apache Iceberg

Learn about the known issues in Iceberg, the impact or changes to the functionality, and the workaround.

CDPD-75667: Querying an Iceberg table with a `TIMESTAMP_LTZ` column can result in data loss

When you query an Iceberg table that has a `TIMESTAMP_LTZ` column, the query could result in data loss.

When creating Iceberg tables from Spark, set the following Spark configuration to avoid creating columns with the `TIMESTAMP_LTZ` type:

```
spark.sql.timestampType=TIMESTAMP_NTZ
```

Apache JIRA: [IMPALA-13484](#)

CDPD-75649: Spark-Iceberg queries fail due to a Java Virtual Machine (JVM) error

While running longevity tests on Spark-Iceberg queries, the query might fail due to the following JVM error - "A fatal error has been detected by the Java Runtime Environment".

Perform the following steps to resolve the issue:

1. From Cloudera Manager, go to Clusters SPARK3 ON YARN Configuration .

2. Search for the "Spark 3 Client Advanced Configuration Snippet (Safety Valve) for spark3-conf/spark-defaults.conf" property and add the following values:

```
spark.driver.extraJavaOptions=-XX:-UseAES
spark.executor.extraJavaOptions=-XX:-UseAES
```

3. Click Save Changes and restart the Spark 3 service for the changes to take effect.

CDPD-72942: Unable to read Iceberg table from Hive after writing data through Apache Flink

If you create an Iceberg table with default values using Hive and insert data into the table through Apache Flink, you cannot then read the Iceberg table from Hive using the Beeline client, and the query fails with the following error:

```
Error while compiling statement: java.io.IOException: java.io.IOException: Cannot create an instance of InputFormat class org.apache.hadoop.mapred.FileInputFormat as specified in mapredWork!
```

The issue persists even after you use the ALTER TABLE statement to set the engine.hive.enabled table property to "true".

None.

Apache JIRA: [HIVE-28525](#)

CDPD-71962: Hive cannot write to a Spark Iceberg table bucketed by date column

If you have used Spark to create an Iceberg table that is bucketed by the "date" column and then try inserting or updating this Iceberg table using Hive, the query fails with the following error:

```
Error: Error while compiling statement: FAILED: RuntimeException org.apache.hadoop.hive.ql.exec.UDFArgumentException: ICEBERG_BUCKET() only takes STRING/CHAR/VARCHAR/BINARY/INT/LONG/DECIMAL/FLOAT/DOUBLE types as first argument, got DATE (state=42000,code=40000)
```

This issue does not occur if the Iceberg table is created through Hive.

None.

Known Issues in Apache Impala

Learn about the known issues in Impala, the impact or changes to the functionality, and the workaround.

IMPALA-532: Impala should tolerate bad locale settings

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

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

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

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

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

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

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

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

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

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

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

IMPALA-3094: Incorrect result due to constant evaluation in query with outer join

An OUTER JOIN query could omit some expected result rows due to a constant such as FALSE in another join clause. For example:

```
explain SELECT 1 FROM alltypestest a1
  INNER JOIN alltypesagg a2 ON a1.smallint_col = a2.year AND fals
e
  RIGHT JOIN alltypes a3 ON a1.year = a3.bigint_col;
+-----+
| Explain String |
+-----+
| Estimated Per-Host Requirements: Memory=1.00KB VCores=1 |
| 00:EMPTYSET |
+-----+
```

CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-59625: Impala shell in RHEL 9 with Python 2 as default does not work

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`.

CDPD-42958: After upgrading the CDH 7.1.9 from CDH 6.x, under certain conditions you cannot insert data into a table

Under the following conditions, after upgrading from CDH 6.x to CDH 7.1.9 you cannot insert data into a table from Impala:

- On CDH 6.x, you created a database with Impala in a user specified HDFS location.
- Using Hive, you then created a table in the database.

Under these conditions, the database and table are stored in the user-specified HDFS directory. After upgrading, the HDFS directory of the table is read-only for Impala. Consequently, from Impala you cannot insert new data into the table because Impala does not have write permission on the HDFS directory.

Workaround: To resolve this issue, use either one of the following workarounds:

- Using the Ranger Web UI, in the policy repository `cm_hdfs`, grant the user 'impala' write permission on the directory where the table resides.
- Enter the following command to grant write permission to user 'impala' on the HDFS directory where the table resides.

```
hdfs dfs -setfacl -m default:user:impala:rwX <HDFS directory>
```

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>/xxxxx
```

Use the [REFRESH/INVALIDATE](#) statements on the affected table to overcome the 'File does not exist' exception.

CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes.

You must use a single Knox node to access Impala UI.

Impala api calls via kinox require configuration if the kinox customized kerberos principal name is a default service user name

To access impala api calls via kinox, if the kinox customized kerberos principal name is a default service user name, then configure "authorized_proxy_user_config" by clicking Clusters->impala->configuration. Include the kinox 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-21828: Multiple permission assignment through grant is not working

None

Problem configuring masking on tables using Ranger

The following Knowledge Base article describes the behavior when we configure masking on tables using Ranger. This configuration works for Hive, but breaks queries in some scenarios for Impala.

For a workaround, see the following Knowledge Base article: [ERROR: "AnalysisException: No matching function with signature: mask\(FLOAT\)" when Impala jobs fail with the following error with signature: mask\(FLOAT\)](#)

IMPALA-11871: INSERT statement does not respect Ranger policies for HDFS

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.

Known Issues in Apache Kafka

Learn about the known issues in Kafka, the impact or changes to the functionality, and the workaround.

Known Issues

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

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

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

1. In Cloudera Manager, go to SMM Configuration Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve)

2. Override bootstrap server URL (hostname:port as set in the listeners for broker) for streams-messaging-manager.yaml.
3. Save your changes.
4. Restart SMM.

The `offsets.topic.replication.factor` property must be less than or equal to the number of live brokers

The `offsets.topic.replication.factor` broker configuration is now enforced upon auto topic creation. Internal auto topic creation will fail with a `GROUP_COORDINATOR_NOT_AVAILABLE` error until the cluster size meets this replication factor requirement.

None

Requests fail when sending to a nonexistent topic with `auto.create.topics.enable` set to true

The first few produce requests fail when sending to a nonexistent topic with `auto.create.topics.enable` set to true.

Increase the number of retries in the producer configuration setting retries.

KAFKA-2561: Performance degradation when SSL Is enabled

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.

CDPD-45183: Kafka Connect active topics might be visible to unauthorised users

The Kafka Connect active topics endpoint (`/connectors/[***CONNECTOR NAME***/topics)` and the Connect Cluster page on the SMM UI disregard the user permissions configured for the Kafka service in Ranger. As a result, all active topics of connectors might become visible to users who do not have permissions to view them. Note that user permission configured for Kafka Connect in Ranger are not affected by this issue and are correctly applied.

None.

RANGER-3809: Idempotent Kafka producer fails to initialize due to an authorization failure

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

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

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

The Debezium Db2 Source connector does not support the evolution (updates) of schemas. In addition, schema change events are not emitted to the schema change topic if there is a change in the schema of a table that is in capture mode. For more information, see [DBZ-4990](#).

None.

CFM-3532: The Stateless NiFi Source, Stateless NiFi Sink, and HDFS Stateless Sink connectors cannot use Snappy compression

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 `libsnapy.so` was copied to the directory you created.
5. Remove the following from `/opt/nativelibs`.

```
libhadoop.a
libhadoop.so
libhadoop.so.1.0.0
```

- 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.

- 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.

- Extract the downloaded archive.
- 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`.

- 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
```

- Change the owner of every entry within `/opt/nativelibs` to `kafka`.

```
chown -h kafka:kafka /opt/nativelibs/*
```

- In Cloudera Manager, go to `Kafka service Configuration`.
- Add the following key-value pair to `Kafka Connect Environment Advanced Configuration Snippet (Safety Valve)`.
 - Key: `LD_LIBRARY_PATH`
 - Value: `/opt/nativelibs`
- Click `Save Changes`.
- Restart the Kafka service.

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

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

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

Unsupported Features

The following Kafka features are not supported in Cloudera Data Platform:

- 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.

- Kafka KRaft in this release of Cloudera Runtime is in technical preview and does not support the following:
 - Deployments with multiple log directories. This includes deployments that use JBOD for storage.
 - Delegation token based authentication.
 - Migrating an already running Kafka service from ZooKeeper to KRaft.
 - Atlas Integration.

Limitations

Collection of Partition Level Metrics May Cause Cloudera Manager's Performance to Degrade

If the Kafka service operates with a large number of partitions, collection of partition level metrics may cause Cloudera Manager's 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 SMM to monitor Kafka or Cruise Control for rebalancing Kafka partitions, be aware that both SMM and Cruise Control rely on partition level metrics. If partition level metric collection is disabled, SMM 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 Kerberos

There are no known issues for Kerberos in Cloudera Runtime 7.3.1.0.

Known Issues in Apache Knox

Learn about the known issues in Knox, the impact or changes to the functionality, and the workaround.

CDPD-76294: Knox service can not be started in a large size Private Cloud Base Cloudera Manager cluster

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 Heapspace 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

Users trying 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 KNOX_GATEWAY_DATA_DIR="/var/lib/knox/gateway/data"; export KNOX_GATEWAY_CONF_DIR="/var/lib/knox/gateway/conf"
3. /opt/cloudera/parcels/CDH/lib/knox/bin/knoxcli.sh create-alias <ALIAS_NAME> <ALIAS_VALUE>
4. 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

If a user makes a concurrent impala-shell connection through Knox, then the connection fails.

Use only one Knox role.

CDPD-60379: During rolling upgrade of Knox service, access fails with 503/500/404/403 error code

The user operation which is performed during the rolling upgrade of Knox might fail with 503/500/404/403 error code.

Retry the user operation.

CDPD-3125: Logging out of Atlas does not manage the external authentication

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

You must use a single Knox node to access Impala UI.

CDPD-22785: Improvements and issues needs to be addressed in convert-topology Knox CLI command

None.

Knox issue with JDK version

jdk-1.8.0_391 is not supported.

Cloudera recommends using Cloudera supported JDKs.

Known Issues in Apache Kudu

Learn about the known issues in Kudu, the impact or changes to the functionality, and the workaround.

- Kudu HMS Sync is disabled and is not yet supported

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
```

KUDU-3619: Major delta compaction for a tablet might fail for particular workloads due to a bug introduced with KUDU-3367

A bug has been introduced with KUDU-3367 functionality. The bug manifests itself when a tablet server's maintenance thread attempts to run a major delta compaction on a tablet where many rows have been deleted, and the attempt fails with an error. To know more about the error message pattern, see KUDU-3619. If that happens, the corresponding tablet might accumulate a lot of updates that cannot be compacted and later garbage collected. In extreme cases, it could lead to running out of disk space when many tablet replicas hosted at the same tablet server hit the issue.

If a tablet server is affected by the issue, messages like the below are present in the tablet server's logs, where <tabletUUID> and <rowsetID> placeholders are populated with corresponding identifiers:

Major delta compaction failed on <tabletUUID>: Corruption: Failed major delta compaction on RowSet(<rowsetID>): No min key found: CFile base data in RowSet(<rowsetID>).

Set the --all_delete_op_delta_file_cnt_for_compaction flag to a very high value (e.g. 1000000) using the Tablet Server Advanced Configuration Snippet (Safety Valve) for gflagfile in the Cloudera Manager UI and restart all the tablet servers in the Kudu cluster.

Apache Jira: [KUDU-3619](#)

Known Issues in Navigator Encrypt

Learn about the known issues in Navigator Encrypt, the impact or changes to the functionality, and the workaround.

Failed to start navencrypt-mount.service

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..

Workaround: Run the command :

```
(sudo) systemctl daemon-reload
```

Known Issues in Apache Oozie

Learn about the known issues in Oozie, the impact or changes to the functionality, and the workaround.

Oozie jobs fail (gracefully) on secure YARN clusters when JobHistory server is down

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.

CDPD-5340: The resourceManager property defined in an Oozie workflow might not work properly if the workflow is submitted through Knox proxy.

An Oozie workflow defined to use the resourceManager property might not work as expected in situations when the workflow is submitted through Knox proxy.

Workaround: Define the jobTracker property with the same value as that of the resourceManager property.

Unsupported Feature

The following Oozie features are currently not supported in Cloudera Data Platform:

- Non-support for Pig action (CDPD-1070)
- Conditional coordinator input logic

Known Issues in Ozone

Learn about the known issues in Ozone, the impact or changes to the functionality, and the workaround.

CDPD-73292: DataNode is supposed to log messages for slow requests to dn-audit.log. However, due to a count error, it will log almost every request.

Set the DataNode configuration `hdds.datanode.slow.op.warning.threshold = 500000000ms`. This configuration change will log only requests that complete in more than 500ms.

OPSAPS-71342: Configuring the hdds.x509.max.duration parameter to 0 or any negative value leads to shutdown of SCM, DN, and OM. This misconfiguration disrupts the entire cluster operations.

To avoid disruption, ensure that the value of `hdds.x509.max.duration` is set to a positive integer greater than 0.

ENGESC-26990: Standalone Ozone deployment in a cluster is not supported.

A cluster with Ozone but without HDFS is not supported. Setting the `fs.defaultFs` parameter to Ozone is not certified yet.

None.

The Prometheus binary is not available in CDP Private Cloud Base 7.1.9 SP1 for the Ubuntu operating system.

You can install Prometheus separately and specify the path to the parent directory, for example `/usr/bin`, in the `prometheus.location` parameter in Ozone.

In CDP Private Cloud Base 7.1.9 SP1, there is a missing block location in the output of Ozone debug chunkinfo command for EC bucket.

None.

In CDP Private Cloud Base 7.1.9 SP1 and CDP Private Cloud Base 7.1.7 SP3, there is an intermittent key put failure after stopping the OM leader.

None.

In CDP Private Cloud 7.1.9 SP1, Recon certificates are not available after running the `ozone admin cert list` command. This issue is present when upgrading from a lower version of CDP to CDP Private Cloud 7.1.9 SP1.

1. Ensure that the Recon user is included in the `hdds.security.client.datanode.container.protocol.acl` configuration in the `hadoop-policy.xml` file under `/var/run/cloudera-scm-agent/process/*-ozone-STORAGE_CONTAINER_MANAGER/ozone-conf`.
2. If the Recon user is not listed, SCM rejects Recon's request for certificates. This step must be completed before proceeding to the next step.
3. Remove the keys folder in the Recon metadata directory and restart the cluster. This action results in new keys and certificates being generated from SCM.

On upgrading the cluster from CDP Private Cloud Base 7.1.7 SP3 to CDP Private Cloud Base 7.1.9 SP1, you cannot create snapshots on a pre-upgrade volume or bucket and PERMISSION_DENIED error is displayed. This is because from CDP Private Cloud Base 7.1.7 SP3, complete Kerberos owner name with domain is used. However, in CDP Private Cloud Base 7.1.9 SP1 the Kerberos owner short name is used.

You must update the Kerberos owner name for the bucket with the user short name and then create snapshot by running the following commands:

- `ozone sh bucket update --user={user_short_name} {volume}/{bucket}`
- `ozone sh snapshot create {volume}/{bucket} {snapshot_name}`

CDPD-69017: An Ozone Manager fails to retrieve the certificates of other OM. OM retrieval of certificates is necessary for delegation token verification. This problem occurs after an OM leader changes until jobs possess a ticket from the previous OM leader.

Copy the certificates of other OM to the certificate directories of all OM.

The [JDK-8292158](#) and [HADOOP-18699](#) bugs affect the following OpenJDK versions:

- OpenJDK 11 (versions lower than 11.0.18)
- OpenJDK 11 (versions lower than 11.0.18-oracle)
- OpenJDK 15 (versions lower than 15.0.10)
- OpenJDK 17 (versions lower than 17.0.6)
- OpenJDK 17 (versions lower than 17.0.6-oracle)
- OpenJDK 19 (versions lower than 19.0.2)

As a result, the Hadoop clients can experience network connection failure under the following conditions:

- The host is capable of supporting AVX-512 instruction sets.
- AVX-512 is enabled in Java Virtual Machine (JVM). This should be enabled by default on AVX-512 capable hosts, equivalent to specifying the JVM argument `-XX:UseAVX=3`
- Hadoop native library (for example, `libhadoop.so`) is not available. So the HDFS client falls back using Hotspot JVM's `aesctr_encrypt` implementation for AES/CTR/NoPadding.
- Client uses an affected JDK.

You must append `-XX:UseAVX=2` to the client JVM args and upgrade to one of the following OpenJDK versions which has the fix:

- OpenJDK 11 release: version 11.0.18 and higher
- OpenJDK 11 release: version 11.0.18-oracle and higher
- OpenJDK 15 release: version 15.0.10 and higher
- OpenJDK 17 release: version 17.0.6 and higher
- OpenJDK 17 release: version 17.0.6-oracle and higher
- OpenJDK 19 release: version 19.0.2 and higher

After upgrading the cluster from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and if Ozone is in the Non-HA environment, an exception message is observed during the finalization of the Ozone upgrade.

During the finalization of the upgrade, a `ClassNotFoundException` message for `org.cloudera.log4j.redactor.RedactorAppender` class may be displayed. The error message can be ignored because the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

CDPD-68951: In CDP Private Cloud Base 7.1.9 CHF2 version and lower, the command `ozone sh key list <bucket_path>` displays the `isFile` flag in a key's metadata as `false` even when the key is a file. This issue is rectified in CDP Private Cloud Base 7.1.9 CHF3. However, the pre-existing (pre-upgrade) key's metadata cannot be changed.

None

When using S3A committer `fs.s3a.committer.name=directory` with `fs.s3a.committer.staging.conflict-mode=replace` to write to FSO buckets, the client fails with the following error.

```
DIRECTORY_NOT_FOUND
org.apache.hadoop.ozone.om.exceptions.OMException:
Failed to find parent directory of xxxxxxxx at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentId(OMFileRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.protocolPB.OzoneManagerRequestHandler.handleWriteRequest
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.runCommand(OzoneManager
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.lambda
$1(OzoneManagerStateMachine.java:363) at
java.base/java.util.concurrent.CompletableFuture
$AsyncSupply.run(CompletableFuture.java:1700) at java.base/
java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1128)
at java.base/java.util.concurrent.ThreadPoolExecutor
$Worker.run(ThreadPoolExecutor.java:628) at java.base/
java.lang.Thread.run(Thread.java:834)
```

This occurs because S3A uses multipart upload to commit job results in a batch. The staging committer's replace mode deletes the target directory before completing Memory Protection Unit (MPU). The problem is that File System Optimization (FSO) does not create intermediate directories during MPU, it does only for regular file/dir/key requests.

Use `fs.s3a.committer.name=magic` for affected versions.

HDDS-9512: Ozone DataNode's new client port conflicts with HDFS DataNode's web port if both Ozone and HDFS DataNode roles are placed on the same host.

You must set `hdds.datanode.client.port` to any unused port. For example, 19864, through the Ozone DataNode safety valve.

OPSAPS-68159: If you did not deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Error when distributing parcel to host, Ignoring non-compliant parcel manifest error is displayed after Cloudera Manager is upgraded to 7.11.3.

If you encounter the error, perform the following steps:

1. Deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.11.3.
2. Restart the cluster with a delay of 10 minutes.
3. Continue to upgrade CDH 7.1.8 to CDH 7.1.9.

OPSAPS-68159: If you did not deactivate the Ozone parcel 718.2.x on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Ozone roles during the CDH 7.1.8 upgrade to CDH 7.1.9.

If you encounter the error, perform the following steps:

1. Deactivate the Ozone parcel 718.2.x.
2. Restart the Ozone service.
3. Perform Finalize Upgrade for the Ozone service.

Step result: The Ozone roles will be displayed in green.

CDPD-60989: The packaging version for Apache Ozone points to the 1.2.0 older version. This is a version string problem and not a packaging issue. The version of the Apache Ozone binary is closest to 1.4.0.

None. This only affects the JAR versioning.

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

OPSAPS-63510: When Ozone Container Balancer is started using Activate Container Balancer from Cloudera Manager, it runs on the Storage Container Manager (SCM) host which is the Ratis leader. However, the link to the Full Log File under Role Log in the Cloudera Manager command window for the Activate Container Balancer command may not link to the leader SCM's logs.

1. Find the leader SCM. Using Cloudera Manager and SCM Web UI: Go to Clusters Ozone **Web UI** . Open any of the Storage Container Manager web UI. In the UI, search for SCM Roles (HA) in the Status section. The leader SCM's hostname is mentioned. Or using Terminal: Log in to any Ozone host and run `ozone admin scm roles`. Note the leader.
2. After finding the leader SCM, search in this leader host's logs for ContainerBalancer related logs.

OPSAPS-67373: Toggling the Enable Ozone S3 Multi-Tenancy feature configuration in the Cloudera Manager Ozone service configuration page affects more service roles than actually needed.

Enabling multi-tenancy only requires restarting the Ozone Managers.

OPSAPS-67757: Hive external tables in Ozone storage cannot be replicated using Hive external table replication policies.

To replicate the Hive external tables' data, consider using DistCp. To replicate the metadata of Hive external tables, consider using HMS Mirror.

Remove bucket recursively using `rb --force` command from AWS S3 cannot work for FSO buckets.

Use the Ozone shell command `ozone sh bucket delete -r [***BUCKET ADDRESS***]`

CDPD-59126: Info log displays "noexec permission on /tmp/liborg_apache_ratis_thirdparty_netty_transport_native_epoll_x86" on client while executing command with noexec on /tmp.

To suppress Info log related to `liborg_apache_ratis_thirdparty_netty_transport_native_epoll_x86` library: Export `OZONE_OPTS` environment variable on the client terminal by running the command `export OZONE_OPTS="-Dorg.apache.ratis.thirdparty.io.netty.native.workdir=/var/tmp $OZONE_OPTS"`

OPSAPS-67650: Ozone uses RocksDB as a library to persist metadata locally.

By default, RocksDB places certain executables in `/tmp`, and thus encounters errors when `/tmp` is mounted with `noexec`.

The workaround is to configure RocksDB to put executables at another location. On a PhatCat node, the steps are:

1. Go to Cloudera Manager UI OZONE Configuration .
2. Find Ozone Service Environment Advanced Configuration Snippet (Safety Valve) and set the following environment variable: `ROCKSDB_SHARED_LIB_DIR=/var/tmp`
3. Restart Ozone.

CDPD-49137: Ozone Manager Kerberos token expires for SCM communication and OM does not log in again.

Sometimes, OM's Kerberos token is not updated and it stops to communicate with SCM. When this occurs, writes start failing.

Restart OM or set the safety valve `hadoop.kerberos.keytab.login.autorenewal.enabled = true`

CDPD-56684: Keys get deleted when you do not have permission on volume

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

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

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.

OPSAPS-67607: Cloudera Manager FirstRun failure at the “Upload YARN MapReduce Framework JARs” step.

If this failure is attributed to the broken symbolic link, /var/lib/hadoop-hdfs/ozone-file-system-hadoop3.jar, it is likely due to the presence of the user hdfs on the node prior to CDP parcel activation. As a result, the Cloudera Manager agent skips the initialization related to HDFS, leading to the non-creation of the /var/lib/hadoop-hdfs directory.

Create the directory “/var/lib/hadoop-hdfs” on all nodes followed by the deactivation and activation of the CDP parcel (deactivate and activate the Ozone parcel instead, in case Ozone parcel is used).

OPSAPS-66501: Currently it is not possible to configure High Availability for SCM roles in Ozone post deployment. You should be able to change the HA configuration through Cloudera Manager, bringing it in line with other services.

At present it requires deleting Ozone and then adding it back with the SCM HA configuration in place and manually cleaning up the Ozone data in between. For more information, read the [KB article](#).

OPSAPS-66500: Currently, it is not possible to enable Kerberos in Ozone after it has been deployed, despite all the required configuration changes being created when the box is checked in the Ozone configurations in Cloudera Manager.

Ozone must be deleted and redeployed with Kerberos enabled. Due to OPSAPS-66499, this requires manual data cleanup in between. For more information, read the [KB article](#).

OPSAPS-66499: When you delete Ozone from a cluster using Cloudera Manager, Ozone data is not cleaned up. This may cause issues when Ozone is redeployed.

You must clean up the data manually. For more information, read the [KB article](#).

CDPD-49027: SCM certificates are not renewed automatically

The certificates that are there to ensure encrypted communication and authentication between Ozone internal services are not renewed automatically for Storage Container Managers.

Certificate revocation

Once these certificates expire, a manual re-bootstrap of the internal Ozone certificates is necessary.

To revoke a certificate, remove the full trust chain to stop trusting a compromised certificate. For this, remove the SCM certificates or any other certificates from the system. During the startup of the

system, new certificates are created and distributed. The old certificates are not trusted anymore as the root CA certificate changes as well.

Procedure to force revoke internal certificates:

1. Stop Ozone service and all of its roles including SCMs
2. Include SCM's certs folders. Note that the Primordial SCM node has two certs folder, one for the root CA and other for the intermediate CA that the node holds. Rest of the SCMs have just one folder for the intermediate CA role that the node serves. The modified command is: `find / -name ozone-metadata 2>/dev/null | while read line; do find $line -name certs; done`
3. Move these certs directories to a backup location
4. Locate the key material and move it to a backup folder. The modified command is: `find / -name ozone-metadata 2>/dev/null | while read line; do find $line -name keys; done`
5. Move these keys directories to a backup location
6. The VERSION file of SCM has to be updated similarly to Ozone Manager's VERSION file. To locate both the SCM and OM VERSION files on the hosts, execute the following command: `find / -name om -o -name scm 2>/dev/null | while read line; do find $line -name VERSION; done | sort | uniq`
7. Backup the version file (just in case you need to restore for any reason)
8. In OM's VERSION file remove the line starting with `omCertSerialId`, in SCM's VERSION file remove the line starting with `scmCertSerialId`.
9. Start the stopped Ozone roles and certificates will be regenerated during startup.

CDPD-35632: The default block level checksum does not work when running distcp from HDFS to Ozone or the other way around, because the two file systems manage underlying blocks very differently.

Use a file level checksum instead. For example, append ``-Ddfs.checksum.combine.mode=COMPOSITE_CRC`` to the distcp command.

CDPD-43942: Requests to modify an Ozone S3 tenant may fail with the error "Timed out acquiring authorizer write lock. Another multi-tenancy request is in-progress." even if another request is not in progress.

Retry the request.

CDPD-22519: HDFS user is unable to run Ozone SCM client CLI.

SCM client CLIs are run using SCM user.

CDPD-34187: This is an usability issue where warnings are displayed on the console while running Ozone fs/CLI commands, which restricts user experience. .

Instead of logging into the user console, you redirect these log messages to a file called `ozone-shell-log4j.properties` which should avoid warnings to the user. Ozone-shell commands used a similar method of directing messages to the LogFile.

CDPD-40594: Ozone admin container create command does not work. The command fails at getCAList for the SCM Client to create a container.

Avoid using the create container command

CDPD-40966: df command on Ozone returns incorrect result.

None

CDPD-34867: Container Balancer might not balance if only Over-Utilized or only Under-Utilized DataNodes are reported. The log line displays the "Container Balancer has identified x Over-Utilized and y Under-Utilized DataNodes that need to be balanced" message where one of x or y will be 0.

Decrease the threshold using `"utilization.threshold"`. This allows the balancer to find non zero number of both over and under utilized nodes.

CDPD-12966: Ozone du -s -h should report correct values with replication information.

None

CDPD-31910: In a non-Ranger deployment, the owner/group are shown based on Kerberos user or sudo user.

For correct owner/group, user needs a Ranger deployment.

CDPD-42691: During the upgrade - all pipelines will be closed when the upgrade is finalized on SCM, temporarily bringing the cluster to a read-only state.

When you execute the finalize command, the cluster will temporarily go into a read-only state.

CDPD-42945: When many EC buckets are created with different EC chunk sizes, it creates pipeline for each chunk size. As a result, large number of pipelines are created in the system.

None

OPSAPS-60721: Ozone SCM Primordial Node ID is a required field which needs to be specified with one of the SCM hostnames during Ozone HA installation. In Cloudera Manager this field is not mandatory during Ozone deployment. This can cause end users to continue further with installation which causes startup to fail in Ozone services.

During Ozone HA installation make sure that Ozone SCM Primordial Node ID is specified with one of the SCM hostname.

HDDS-4209: S3A Filesystem does not work with Ozone S3 in file system compact mode. When you create a directory, the S3A filesystem creates an empty file. When the ozone.om.enable.filesystem.paths parameter is enabled, the `hdfs dfs -mkdir -p s3a://b12345/d11/d12` command runs successfully. However, running the `hdfs dfs -put /tmp/file1 s3a://b12345/d11/d12/file1` command fails with an error: `ERROR org.apache.hadoop.ozone.om.request.key.OMKeyCreateRequest: Key creation failed.`

The HDDS-4209 Jira fixes the file system semantics and management in Ozone. On top of the flat name structure, which is Pure Object store, as a workaround the Hierarchical namespace structure is added. This ensures S3A compatibility with Ozone.

CDPD-41539: The "No such file or directory" imessage is displayed when EC file is read using older ofs client.

You must upgrade the client before trying to read the key: `vol1/ecbuck1/1GB_ec`.

CDPD-42832: With this issue, any long running setup or a production server results in data corruption due to inconsistency issues. This may result in major issues with the existing Legacy layout type.

FSO provides atomicity and consistency guarantees for the path (dir or file) rename/delete operations irrespective of the large sub-dirs/files contained in it. These capabilities help to make the long running test more consistent without any failures so far. Recommendation is to run bigdata HCFS workloads using the FSO bucket layout types.

OPSAPS-63999: In the newly installed cluster, the Finish upgrade option is clickable.

None

CDPD-45932: Investigate impersonation with "is admin" in Ozone WebUIs /logLevel servlet endpoint

In a secure Kerberized cluster, due to an impersonation issue, changing log levels using Knox on the corresponding endpoint of the WebUI does not work. Note that this is only true, when the WebUI is accessed using Knox. Other means of changing log levels in Ozone services are not affected by this problem.

None.

CDPD-74201: ozone sh key list prints the `--all` and `--length` options twice. This is a listing issue.

None.

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-74045: Destination cluster's Ozone services stops running when 30 replication policies or more are run concurrently with large number of files.

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.

None.

CDPD-74483: Ozone replication from source to destination fails when `ozone.client.bytes.per.checksum` is set to 16KB.

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.

None.

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.

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.

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.

You must wait for SCM to come out of samemode or exit from safemode through CLI options.

CDPD-74257: For the 7.3.1 release, the Heatmap feature is disabled and removed from the Recon UI and the Solr health check is removed from the Recon Overview page.

If the APIs are called, heatmap APIs will block threads if the `javax.security.auth.useSubjectCredsOnly` property is set to true.

To check if `javax.security.auth.useSubjectCredsOnly` is set to true, you can run the `sudo -u hdfs /usr/java/jdk1.8.0_232-cloudera/bin/jcmd <PID> VM.system_properties | grep -i subject` command where PID must be replaced with Recon process id.

To identify the PID, run `ps -ef` and `grep` for Recon to find the Java process running Recon.



Important: If `javax.security.auth.useSubjectCredsOnly` property is set to true, you must restart the cluster.



Warning: After restarting the cluster, any background Ozone jobs running might be affected including data corruption in case read or write operations were performed. You must restart only when:

- You are aware of the flag set
- You must ensure no critical background jobs are running

Cloudera does not recommend you call any Heatmap API on the CDP Private Cloud Base 7.3.1 cluster to prevent blocking.

CDPD-75958/CDPD-75954: The ozone debug ldb command and ozone auditparser fails with java.lang.UnsatisfiedLinkError.

For information on workaround, see the [Changing /tmp directory for CLI tools](#) documentation.

OPSAPS-72144: During the Finalize Upgrade process, the SCM user and ozone.keytab is used by default. This causes access issues in using a custom Kerberos principal. For example, scmFoo, results in an Access denied error, as the SCM superuser privilege is required.

Finalize Upgrade command fails with an access denied error for scm/host@DOMAIN. Logs indicate the use of kinit with the default SCM user rather than the custom Kerberos principal.

You must add the Custom Kerberos Principal to Ozone Administrators.

1. Log in to Cloudera Manager UI.
2. Navigate to Clusters.
3. Select the Ozone service.
4. Go to Configurations.
5. Search for ozone.administrators.
6. Add the short form of the custom Kerberos principal. For example, scmFoo without the domain suffix
7. Click Save Changes.
8. Restart the Ozone service.
9. Run the Finalize Upgrade command.

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

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

Known Issues in Apache Parquet

There are no known issues for Parquet in Cloudera Runtime 7.3.1.0.

Known Issues in Apache Phoenix

There are no known issues for Phoenix in Cloudera Runtime 7.3.1.

Known Issues in Apache Ranger

Learn about the known issues in Ranger, the impact or changes to the functionality, and the workaround.

Ranger Tagsync does not support Ozone OFS paths / O3FS recursive feature not supported (Tag based access control behavior may not be as expected for ozone)

There is no support for OFS path/O3FS recursive feature in 7.3.1. If upgrading from 7.1.9 SP1 CHF3 or higher, there will be a regression.

Wait for the next SP/CHF release 7.3.1 before upgrading.

CDPD-75532: Remove self node from the resourceTrie only if it has no children, no evaluators and no wildcard-evaluators

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.

Restart the plugin service whose policy is not being evaluated.

CDPD-68806: The Revoke operation for users belonging to a group or role permission does not function as expected

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

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-67238: Multiple Columns Revoke not generating policies with correct number of columns

As an example, when "revoke select(col1, col2,col3) on table demo.test from role Role3;" is done, the generated policy does not revoke the columns. Currently the revoke statement is only revoking if there is only one column.

None.

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-58704: hadoop roll key/key delete command shows operation failed error when one KMS host is down, even when operation succeeds

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

None.

CDPD-56741: Improvement in log message when jwtauth not used

None.

CDPD-56738: Ranger RMS showing FileNotFoundException: /usr/share/java/oraclepki.jar in Oracle 19 setup

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-55107: Not able to search using multiple user filter in access audit tab

If you were using multiple user search filters in Audit Access Tab on Ranger Admin UI, after upgrading to CDP-7.1.9 that would not be supported. You can continue to search users with a single search filter.

None.

CDPD-48975: Ranger KMS KTS to KMS DB migration : keys with the same name but different case are not migrated

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

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": []}}]

Expectation is to return all the classifications. But the response has only "classification". Happens similarly for entity-label, entity-business-metadata.

None.

CDPD-40734: User allowed to insert data into a hive table when there is a deny policy on a table column

A user is allowed to enter data into a table even if there is a deny policy present on one of the table columns.

The user is able to insert data into the table.

None.

Known Issues in Ranger KMS

Learn about the known issues in Ranger KMS, the impact or changes to the functionality, and the workaround.

None.

Known Issues in Schema Registry

Learn about the known issues in Schema Registry, the impact or changes to the functionality, and the workaround.

CDPD-40380: Authorization checking issue when Kerberos is disabled

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

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

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

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.

HBase indexer does not load netty and snappy libraries

The HBase indexer loads the netty and snappy libraries and these libraries are necessary for the Key-Value Indexer to work. However, the Key-Value Indexer cannot use these libraries if `/tmp` is mounted with the `noexec` property. To address this issue, you have to manually specify another directory instead of the default `/tmp`.

Perform the following steps to resolve this issue:

1. Go to the Key-Value Store Indexer service Configuration .
2. Search for the Key-Value Store Indexer Service Environment Advanced Configuration Snippet (Safety Valve) property.
3. If the `HBASE_INDEXER_OPTS` key is already present in the configuration, append the following value else add the following key and value:

```
Name: HBASE_INDEXER_OPTS
Value: -Dorg.apache.hbase.thirdparty.io.netty.native.workdir=/var/hbase-solr/netty-workdir -Dorg.xerial.snappy.tempdir=/var/hbase-solr/snappy-tempdir
```



Note:

- If the `/var/hbase-solr/netty-workdir` and `/var/hbase-solr/snappy-tempdir` file system locations do not exist, create the directories and ensure that the "hbase" user has permissions to write into these directories.
- Run the `chown` command on the directories. For example,

```
chown -R hbase:hbase /var/hbase-solr/netty-workdir
```

4. Restart the Key-Value Store Indexer service by clicking Key-Value Store Indexer service Actions Restart .

HBase Indexer does not work with JDK 17

Depending on the Cloudera Manager version used with CDP, HBase Indexer (KS Indexer) may have compatibility issues with JDK 17.

You have the following options to fix this issue:

- Upgrade Cloudera Manager to version 7.11.3 or higher.
- If upgrading Cloudera Manager is not an option, you can manually add the following to HBase Indexer Java options in Cloudera Manager:

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

Splitshard operation fails after CDH 6 to CDP upgrade

Collections are not reindexed during an upgrade from CDH 6 to CDP 7 because Lucene 8 (CDP) can read Lucene 7 (CDH 6) indexes.

If you try to execute a SPLITSHARD operation against such a collection, it fails with a similar error message:

```
o.a.s.h.a.SplitOp ERROR executing split: => java.lang.IllegalArg
umentException: Cannot merge a segment t
hat has been created with major version 7 into this index which
has been created by major version 8
    at org.apache.lucene.index.IndexWriter.validateMergeRea
der(IndexWriter.java:3044)
java.lang.IllegalArgumentException: Cannot merge a segment that h
as been created with major version 7 into this index which has b
een created by major version 8
    at org.apache.lucene.index.IndexWriter.validateMergeReade
r(IndexWriter.java:3044) ~[lucene-core-8.11.2.7.1.9.3-2.jar:8.11
.2.7.1.9.3-2 a6ff93f9665115dffbdad0ad7f222fd1978d495d - jenkins -
2023-12-02 00:05:23]
    at org.apache.lucene.index.IndexWriter.addIndexes(IndexWr
iter.java:3110) ~[lucene-core-8.11.2.7.1.9.3-2.jar:8.11.2.7.1.9.
3-2 a6ff93f9665115dffbdad0ad7f222fd1978d495d - jenkins - 2023-12
-02 00:05:23]
    at org.apache.solr.update.SolrIndexSplitter.doSplit(So
lrIndexSplitter.java:318) ~[solr-core-8.11.2.7.1.9.3-2.jar:8.11.
2.7.1.9.3-2 a6ff93f9665115dffbdad0ad7f222fd1978d495d - jenkins -
2023-12-02 00:16:28]
    at org.apache.solr.update.SolrIndexSplitter.split(Solr
IndexSplitter.java:184) ~[solr-core-8.11.2.7.1.9.3-2.jar:8.11.2.
7.1.9.3-2 a6ff93f9665115dffbdad0ad7f222fd1978d495d - jenkins - 2
023-12-02 00:16:28]
    at org.apache.solr.update.DirectUpdateHandler2.split(Dir
ectUpdateHandler2.java:922) ~[solr-core-8.11.2.7.1.9.3-2.jar:8.1
1.2.7.1.9.3-2 a6ff93f9665115dffbdad0ad7f222fd1978d495d - jenkins
- 2023-12-02 00:16:28]
    at org.apache.solr.handler.admin.SplitOp.execute(SplitOp
.java:165) ~[solr-core-8.11.2.7.1.9.3-2.jar:8.11.2.7.1.9.3-2 a6f
f93f9665115dffbdad0ad7f222fd1978d495d - jenkins - 2023-12-02 00:
16:28]
    at org.apache.solr.handler.admin.CoreAdminOperation.execu
te(CoreAdminOperation.java:367) ~[solr-core-8.11.2.7.1.9.3-2.jar
:8.11.2.7.1.9.3-2 a6ff93f9665115dffbdad0ad7f222fd1978d495d - jen
kins - 2023-12-02 00:16:28]
```

This happens because the segment created using a Lucene 7 index cannot be merged into a Lucene 8 index.

Drop the entire collection, delete the data in HDFS and recreate the collection with Solr 8 configs.

Changing the default value of Client Connection Registry HBase configuration parameter causes HBase MRIT job to fail

If the value of the HBase configuration property Client Connection Registry is changed from the default ZooKeeper Quorum to Master Registry then the Yarn job started by HBase MRIT fails with a similar error message:

```
Caused by: org.apache.hadoop.hbase.exceptions.MasterRegistryFetc
hException: Exception making rpc to masters [quasar-bmyccr-2.qua
sar-bmyccr.root.hwx.site,22001,-1]
    at org.apache.hadoop.hbase.client.MasterRegistry.lambda$g
roupCall$1(MasterRegistry.java:244)
    at org.apache.hadoop.hbase.util.FutureUtils.lambda$addLi
stener$0(FutureUtils.java:68)
```

```

        at java.util.concurrent.CompletableFuture.uniWhenComple
ete(CompletableFuture.java:774)
        at java.util.concurrent.CompletableFuture.uniWhenComple
eStage(CompletableFuture.java:792)
        at java.util.concurrent.CompletableFuture.whenComplete(Co
mpletableFuture.java:2153)
        at org.apache.hadoop.hbase.util.FutureUtils.addListener(F
utureUtils.java:61)
        at org.apache.hadoop.hbase.client.MasterRegistry.groupCa
ll(MasterRegistry.java:228)
        at org.apache.hadoop.hbase.client.MasterRegistry.call(Ma
sterRegistry.java:265)
        at org.apache.hadoop.hbase.client.MasterRegistry.getMetaR
egionLocations(MasterRegistry.java:282)
        at org.apache.hadoop.hbase.client.ConnectionImplementati
on.locateMeta(ConnectionImplementation.java:900)
        at org.apache.hadoop.hbase.client.ConnectionImplementat
ion.locateRegion(ConnectionImplementation.java:867)
        at org.apache.hadoop.hbase.client.ConnectionImplementati
on.relocateRegion(ConnectionImplementation.java:850)
        at org.apache.hadoop.hbase.client.ConnectionImplementat
ion.locateRegionInMeta(ConnectionImplementation.java:981)
        at org.apache.hadoop.hbase.client.ConnectionImplementa
tion.locateRegion(ConnectionImplementation.java:870)
        at org.apache.hadoop.hbase.client.RpcRetryingCallerWith
ReadReplicas.getRegionLocations(RpcRetryingCallerWithReadReplica
s.java:319)
        ... 21 more
Caused by: org.apache.hadoop.hbase.client.RetriesExhaustedExcept
ion: Failed contacting masters after 1 attempts.
Exceptions:
java.io.IOException: Call to address=quasar-bmyccr-2.quasar-bmy
ccr.root.hwx.site/172.27.19.4:22001 failed on local exception: j
ava.io.IOException: java.lang.RuntimeException: Found no valid a
uthentication method from options
        at org.apache.hadoop.hbase.client.MasterRegistry.lambda
$groupCall$1(MasterRegistry.java:243)
        ... 35 more

```

Add the following line to the MRIT command line:

```
-D 'hbase.client.registry.impl=org.apache.hadoop.hbase.client.ZK
ConnectionRegistry'
```

Solr does not support rolling upgrade to release 7.1.9 or lower

Solr supports rolling upgrades from release 7.1.9 and higher. Upgrading from a lower version means that all the Solr Server instances are shut down, parcels upgraded and activated and then the Solr Servers are started again. This causes a service interruption of several minutes, the actual value depending on cluster size.

Services like Atlas and Ranger that depend on Solr, may face issues because of this service interruption.

None.

Unable to see single valued and multivalued empty string values when querying collections after upgrade to CDP

After upgrading from CDH or HDP to CDP, you are not able to see single valued and multi Valued empty string values in CDP.

This behavior in CDP is due to the remove-blank processor present in solrconfig.xml in Solr 8.

Remove the remove-blank processor from solrconfig.xml.

Cannot create multiple heap dump files because of file name error

Heap dump generation fails with a similar error message:

```
java.lang.OutOfMemoryError: Java heap space
Dumping heap to /data/tmp/solr_solr-SOLR_SERVER-fc9dacc265fabfc5
00b92112712505e3_pid{{PID}}.hprof ...
Unable to create /data/tmp/solr_solr-SOLR_SERVER-fc9dacc265fab
fc500b92112712505e3_pid{{PID}}.hprof: File exists
```

The cause of the problem is that {{PID}} does not get substituted during dump file creation with an actual process ID and because of that, a generic file name is generated. This causes the next dump file creation to fail, as the existing file with the same name cannot be overwritten.

You need to manually delete the existing dump file.

Solr coreAdmin status throws Null Pointer Exception

You get a Null Pointer Exception with a similar stacktrace:

```
Caused by: java.lang.NullPointerException
    at org.apache.solr.core.SolrCore.getInstancePath(SolrCore.
java:333)
    at org.apache.solr.handler.admin.CoreAdminOperation.getCor
eStatus(CoreAdminOperation.java:324)
    at org.apache.solr.handler.admin.StatusOp.execute(StatusOp.
java:46)
    at org.apache.solr.handler.admin.CoreAdminOperation.execute
(CoreAdminOperation.java:362)
```

This is caused by an error in handling solr admin core STATUS after collections are rebuilt.

Restart the Solr server.

Applications fail because of mixed authentication methods within dependency chain of services

Using different types of authentication methods within a dependency chain, for example, configuring your indexer tool to authenticate using Kerberos and configuring your Solr Server to use LDAP for authentication may cause your application to time out and eventually fail.

Make sure that all services in a dependency chain use the same type of authentication.

API calls fail with error when used with alias, but work with collection name

API calls fail with a similar error message when used with an alias, but they work when made using the collection name:

```
[    ] o.a.h.s.t.d.w.DelegationTokenAuthenticationFilter Authenti
cation exception: User: xyz@something.example.com is not allowed
to impersonate xyz@something.example.com
[c:RTOTagMetaOdd s:shard3 r:core_node11 x:RTOTagMetaOdd_shar
d3_replica_n8] o.a.h.s.t.d.w.DelegationTokenAuthenticationFilter
Authentication exception: User: xyz@something.example.com is not
allowed to impersonate xyz@something.example.com
```

Make sure there is a replica of the collection on every host.

CrunchIndexerTool does not work out of the box if /tmp is mounted noexec mode

When you try to run CrunchIndexerTool with the /tmp directory mounted in noexec mode, It throws a snappy-related error.

Create a separate directory for snappy temp files which is mounted with EXEC privileges and set this directory as the value of the org.xerial.snappy.tmpdir java property as a driver java option.

For example:

```
export myDriverJarDir=/opt/cloudera/parcels/CDH//lib/solr/contrib/crunch;export myDependencyJarDir=/opt/cloudera/parcels/CDH//lib/search/lib/search-crunch;export myDriverJar=$(find $myDriverJarDir -maxdepth 1 -name 'search-crunch-*.jar' ! -name '*-job.jar' ! -name '*-sources.jar');export myDependencyJarFiles=$(find $myDependencyJarDir -name '*.jar' | sort | tr '\n' ',' | head -c -1);export myDependencyJarPaths=$(find $myDependencyJarDir -name '*.jar' | sort | tr '\n' ':' | head -c -1);export HADOOP_CONF_DIR=;spark-submit --master local --deploy-mode client --driver-library-path /opt/cloudera/parcels/CDH//lib/hadoop/lib/native/ --jars $myDependencyJarFiles --driver-java-options '-Dorg.xerial.snappy.tmpdir=/home/systest/tmp' --class org.apache.solr.crunch.CrunchIndexerTool $myDriverJar --input-file-format=avroParquet --input-file-reader-schema search-parquetfile/parquet-schema.avsc --morphline-file /tmp/mrTestBase.conf --pipeline-type spark --chatty hdfs://[***HOSTNAME***]:8020/tmp/parquetfileparsertest-input
```

Mergeindex operation with --go-live fails after CDH 6 to CDP upgrade

During an upgrade from CDH6 to CDP, collections are not reindexed because Lucene 8 (CDP) can read Lucene 7 (CDH6) indexes.

If you try to execute MapReduceIndexerTool (MRIT) or HBase Indexer MRIT with --go-live against such a collection, you get a similar error message:

```
Caused by: java.lang.IllegalArgumentException: Cannot merge a segment that has been created with major version 8 into this index which has been created by major version 7
    at org.apache.lucene.index.IndexWriter.validateMergeReader(IndexWriter.java:2894)
    at org.apache.lucene.index.IndexWriter.addIndexes(IndexWriter.java:2960)
    at org.apache.solr.update.DirectUpdateHandler2.mergeIndexes(DirectUpdateHandler2.java:570)
    at org.apache.solr.update.processor.RunUpdateProcessor.processMergeIndexes(RunUpdateProcessorFactory.java:95)
    at org.apache.solr.update.processor.UpdateRequestProcessor.processMergeIndexes(UpdateRequestProcessor.java:63)
```

This happens because CDP MRIT and HBase indexer use Solr 8 as embedded Solr, which creates a Lucene 8 index. It cannot be merged (using MERGEINDEXES) into an older Lucene 7 index.

In the case of MRIT the only way to move past this issue is to drop the entire collection, delete the data in HDFS and recreate the collection with Solr 8 configs.

For HBase Indexer MRIT an alternative workaround is setting the number of reducers to 0 (--reducers 0) because in this case documents are sent directly from the mapper tasks to live Solr servers instead of using MERGEINDEXES.

Apache Tika upgrade may break morphlines indexing

The upgrade of Apache Tika from 1.27 to 2.3.0 brought potentially breaking changes for morphlines indexing. Duplicate/triplicate keys names were removed and certain parser class names were changed (For example, org.apache.tika.parser.jpeg.JpegParser changed to org.apache.tika.parser.image.JpegParser).

To avoid morphline commands failing after the upgrade, do the following:

- Check if key name changes affect your morphlines. For more information, see *Removed duplicate/triplicate keys* in [Migrating to Tika 2.0.0](#).

- Check if the name of any parser you use has changed. For more information, see the Apache Tika [API documentation](#).

Update your morphlines if necessary.

CDPD-28006: Solr access via Knox fails with impersonation error though auth_to_local and proxy user configs are set

Currently the names of system users which are impersonating users with Solr should match with the names of their respective Kerberos principals.

If, for some reason, this is not feasible, you must add the user name you want to associate with the custom Kerberos principal to Solr configuration via the Solr Service Environment Advanced Configuration Snippet (Safety Valve) environment variable in Cloudera Manager.

For more information, see [Configuring custom Kerberos principals and custom system users](#).

CDH-77598: Indexing fails with socketTimeout

Starting from CDH 6.0, the HTTP client library used by Solr has a default socket timeout of 10 minutes. Because of this, if a single request sent from an indexer executor to Solr takes more than 10 minutes to be serviced, the indexing process fails with a timeout error.

This timeout has been raised to 24 hours. Nevertheless, there still may be use cases where even this extended timeout period proves insufficient.

If your MapreduceIndexerTool or HBaseMapreduceIndexerTool batch indexing jobs fail with a timeout error during the go-live (Live merge, MERGEINDEXES) phase (This means the merge takes longer than 24 hours).

Use the --go-live-timeout option where the timeout can be specified in milliseconds.

CDPD-12450: CrunchIndexerTool Indexing fails with socketTimeout

The http client library uses a socket timeout of 10 minutes. The Spark Crunch Indexer does not override this value, and in case a single batch takes more than 10 minutes, the entire indexing job fails. This can happen especially if the morphlines contain DeleteByQuery requests.

Try the following workarounds:

- Check the batch size of your indexing job. Sending too large batches to Solr might increase the time needed on the Solr server to process the incoming batch.
- If your indexing job uses deleteByQuery requests, consider using deleteById wherever possible as deleteByQuery involves a complex locking mechanism on the Solr side which makes processing the requests slower.
- Check the number of executors for your Spark Crunch Indexer job. Too many executors can overload the Solr service. You can configure the number of executors by using the --mappers parameter
- Check that your Solr installation is correctly sized to accommodate the indexing load, making sure that the number of Solr servers and the number of shards in your target collection are adequate.
- The socket timeout for the connection can be configured in the morphline file. Add the solrClientSocketTimeout parameter to the solrLocator command

Example

```
SOLR_LOCATOR :
{
  collection : test_collection
  zkHost : "zookeeper1.example.corp:2181/solr"
  # 10 minutes in milliseconds
  solrClientSocketTimeout: 600000
  # Max number of documents to pass per RPC from morphline to
  Solr Server
  # batchSize : 10000
```

```
}
```

CDPD-29289: HBaseMapReduceIndexerTool fails with socketTimeout

The http client library uses a socket timeout of 10 minutes. The HBase Indexer does not override this value, and in case a single batch takes more than 10 minutes, the entire indexing job fails.

You can overwrite the default 600000 millisecond (10 minute) socket timeout in HBase indexer using the `--solr-client-socket-timeout` optional argument for the direct writing mode (when the value of the `--reducers` optional argument is set to 0 and mappers directly send the data to the live Solr).

CDPD-20577: Splitshard operation on HDFS index checks local filesystem and fails

When performing a shard split on an index that is stored on HDFS, `SplitShardCmd` still evaluates free disk space on the local file system of the server where Solr is installed. This may cause the command to fail, perceiving that there is no adequate disk space to perform the shard split.

Run the following command to skip the check for sufficient disk space altogether:

- On nonsecure clusters:

```
curl 'http://$[***SOLR_SERVER_HOSTNAME***]:8983/solr/admin/collections?action=SPLITSHARD&collection=[***COLLECTION_NAME***]&shard=[***SHARD_TO_SPLIT***]&skipFreeSpaceCheck=true'
```

- On secure clusters:

```
curl -k -u : --negotiate 'http://$[***SOLR_SERVER_HOSTNAME***]:8985/solr/admin/collections?action=SPLITSHARD&collection=[***COLLECTION_NAME***]&shard=[***SHARD_TO_SPLIT***]&skipFreeSpaceCheck=true'
```

Replace `[***SOLR_SERVER_HOSTNAME***]` with a valid Solr server hostname, `[***COLLECTION_NAME***]` with the collection name, and `[***SHARD_TO_SPLIT***]` with the ID of the to split.

To verify that the command executed successfully, check overseer logs for a similar entry:

```
2021-02-02 12:43:23.743 INFO (OverseerThreadFactory-9-thread-5-processing-n:myhost.example.com:8983_solr) [c:example s:shard1] o.a.s.c.a.c.SplitShardCmd Skipping check for sufficient disk space
```

CDH-22190: CrunchIndexerTool which includes Spark indexer requires specific input file format specifications

If the `--input-file-format` option is specified with `CrunchIndexerTool`, then its argument must be text, avro, or avroParquet, rather than a fully qualified class name.

None

CDH-26856: Field value class guessing and Automatic schema field addition are not supported with the MapReduceIndexerTool nor with the HBaseMapReduceIndexerTool

The `MapReduceIndexerTool` and the `HBaseMapReduceIndexerTool` can be used with a Managed Schema created via NRT indexing of documents or via the Solr Schema API. However, neither tool supports adding fields automatically to the schema during ingest.

Define the schema before running the `MapReduceIndexerTool` or `HBaseMapReduceIndexerTool`. In non-schemaless mode, define in the schema using the `schema.xml` file. In schemaless mode, either define the schema using the Solr Schema API or index sample documents using NRT indexing before invoking the tools. In either case, Cloudera recommends that you verify that the schema is what you expect, using the List Fields API command.

Users with insufficient Solr permissions may encounter a blank Solr Web Admin UI

Users who are not authorized to use the Solr Admin UI are not given a page explaining that access is denied to them, instead they receive a blank Admin UI with no information.

None

CDH-15441: Using MapReduceIndexerTool or HBaseMapReduceIndexerTool multiple times may produce duplicate entries in a collection

Repeatedly running the MapReduceIndexerTool on the same set of input files can result in duplicate entries in the Solr collection. This occurs because the tool can only insert documents and cannot update or delete existing Solr documents. This issue does not apply to the HBaseMapReduceIndexerTool unless it is run with more than zero reducers.

To avoid this issue, use HBaseMapReduceIndexerTool with zero reducers.



Note: This workaround is only valid for HBaseMapReduceIndexerTool. There is no workaround for MapReduceIndexerTool

CDH-58694: Deleting collections might fail if hosts are unavailable

It is possible to delete a collection when hosts that host some of the collection are unavailable. After such a deletion, if the previously unavailable hosts are brought back online, the deleted collection may be restored.

Ensure all hosts are online before deleting collections.

CDPD-13923: Every Configset is Untrusted Without Kerberos

Solr 8 introduces the concept of ‘[untrusted configset](#)’, denoting configsets that were uploaded without authentication. Collections created with an untrusted configset will not initialize if <lib> directives are used in the configset.

Select one of the following options if you would like to use untrusted configsets with <lib> directives:

- If the configset contains external libraries, but you do not want to use them, simply upload the configsets after deleting the <lib> directives.
- If the configset contains external libraries, and you want to use them, choose one from the following options:
 - Secure your cluster before reuploading the configset.
 - Add the libraries to Solr’s classpath, then reupload the configset without the <lib> directives.

CDPD-71422: Solr went into an unhealthy state after the data lake upgrade

After the Data Lake upgrade to the 7.3.1.0 version, the Solr service becomes unhealthy for the public cloud environments (AWS, Azure, and GCP). This is an intermittent issue.

Manually restart the Solr service in the Data Lake after an upgrade.

Unsupported features

The following Solr features are currently not supported in Cloudera Data Platform:

- Panel with security info in admin UI's dashboard
- Incremental backup mode
- Schema Designer UI
- Package Management System
- HTTP/2
- Solr SQL/JDBC
- Graph Traversal
- Cross Data Center Replication (CDCR)
- SolrCloud Autoscaling
- HDFS Federation

- Saving search results
- Solr contrib modules

(Spark, MapReduce, and Lily HBase indexers are not contrib modules but part of Cloudera's distribution of Solr itself, therefore they are supported)

Limitations

Enabling blockcache writing may result in unusable indexes

It is possible to create indexes with `solr.hdfs.blockcache.write.enabled` set to true. Such indexes may appear corrupt to readers, and reading these indexes may irrecoverably corrupt them. Because of this, blockcache writing is disabled by default.

Default Solr core names cannot be changed

Although it is technically possible to give user-defined Solr core names during core creation, it is to be avoided in the context of Cloudera's distribution of Apache Solr. Cloudera Manager expects core names in the default "collection_shardX_replicaY" format. Altering core names results in Cloudera Manager being unable to fetch Solr metrics for the given core and this may corrupt data collection for co-located core, or even shard, and server level charts.

Lucene index handling limitation

The Lucene index can only be upgraded by one major version. Solr 8 will not open an index that was created with Solr 6 or earlier. Because of this, you need to reindex collections that were created with Solr 6 or earlier.

Known Issues in Apache Spark

Learn about the known issues in Spark, the impact or changes to the functionality, and the workaround.

Spark 3: RAPIDS Accelerator is not available

The RAPIDS Accelerator for Apache Spark is currently not available in Cloudera Private Cloud 7.3.1.0.

Workaround: None

The CHAR(n) type handled inconsistently, depending on whether the table is partitioned or not.

In upstream Spark 3 the `spark.sql.legacy.charVarcharAsString` configuration was introduced, but it does not solve all incompatibilities with Spark 2.

Workaround: 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 for Apache Sqoop

Learn about the known issues in Apache Sqoop, the impact or changes to the functionality, and the workaround.

CDPD-54770: Unable to read Sqoop metastore created by an older HSQLDB version

If you have upgraded to CDP PvC Base 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 CDP PvC Base 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

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

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.

Workaround: Prepend column names in Parquet tables with one or more letters or underscore characters.

Apache JIRA: None

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

Learn about the known issues in Streams Messaging Manager (SMM), the impact or changes to the functionality, and the workaround.

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

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

View the SMM UI logs on the host.

CDPD-39313: Some numbers are not rendered properly in SMM UI

Very large numbers can be imprecisely represented on the UI. For example, bytes larger than 8 petabytes would lose precision.

None.

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

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

SMM 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 SMM 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 SMM.

CDPD-45183: Kafka Connect active topics might be visible to unauthorised users

The Kafka Connect active topics endpoint (/connectors/[***CONNECTOR NAME***/topics) and the Connect Cluster page on the SMM UI disregard the user permissions configured for the Kafka service in Ranger. As a result, all active topics of connectors might become visible to users who do not have permissions to view them. Note that user permission configured for Kafka Connect in Ranger are not affected by this issue and are correctly applied.

None.

Limitations**CDPD-36422: 1MB flow.snapshot freezes Safari**

While importing large connector configurations, flow.snapshots reduces the usability of the Streams Messaging Manager when using Safari browser.

Use a different browser (Chrome/Firefox/Edge).

Known Issues in Streams Replication Manager

Learn about the known issues in Streams Replication Manager (SRM), the impact or changes to the functionality, and the workaround.

Known Issues**CDPD-22089: SRM does not sync re-created source topics until the offsets have caught up with target topic**

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

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 SMM UI. However, its Partitions and Consumer Groups will be 0, its Throughput, Replication Latency and Checkpoint Latency will show N/A.

None

CDPD-30275: SRM may automatically re-create deleted topics on target clusters

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 SRM is no longer replicating the topic.
3. Delete the remote topic in the target cluster.

Limitations

SRM cannot replicate Ranger authorization policies to or from Kafka clusters

Due to a limitation in the Kafka-Ranger plugin, SRM cannot replicate Ranger policies to or from clusters that are configured to use Ranger for authorization. If you are using SRM to replicate data to or from a cluster that uses Ranger, disable authorization policy synchronization in SRM. This can be achieved by clearing the Sync Topic Acls Enabled (sync.topic.acls.enabled) checkbox.

SRM cannot ensure the exactly-once semantics of transactional source topics

SRM data replication uses at-least-once guarantees, and as a result cannot ensure the exactly-once semantics (EOS) of transactional topics in the backup/target cluster.



Note: Even though EOS is not guaranteed, you can still replicate the data of a transactional source, but you must set isolation.level to read_committed for SRM's internal consumers. This can be done by adding `[***SOURCE CLUSTER ALIAS***]->[***TARGET CLUSTER ALIAS***].consumer.isolation.level=read_committed` to the Streams Replication Manager's Replication Configs SRM service property in Cloudera Manager. The isolation.level property can be set on a global connector or replication level. For example:

```
#Global connector level
connectors.consumer.isolation.level=read_committed
#Replication level
uswest->useast.consumer.isolation.level=read_committed
```

SRM checkpointing is not supported for transactional source topics

SRM does not correctly translate checkpoints (committed consumer group offsets) for transactional topics. Checkpointing assumes that the offset mapping function is always increasing, but with transactional source topics this is violated. Transactional topics have control messages in them, which take up an offset in the log, but they are never returned on the consumer API. This causes the mappings to decrease, causing issues in the checkpointing feature. As a result of this limitation, failover operations for transactional topics is not possible.

Known Issues in YARN, YARN Queue Manager and MapReduce

Learn about the known issues in YARN, YARN Queue Manager and MapReduce, the impact or changes to the functionality, and the workaround.

Known Issues

COMPX-14820: Delete Queue and its Children throws "Queue capacity was reduced to zero, but failed to delete queue."

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.

COMPX-13177: QueueManager webapp requests fail with 'HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA'

Products:

- Cloudera Manager for CDP Private
- Cloud Base Cloudera Manager for CDP Public Cloud

Context:

- Centos 7.8 and Redhat 7.8 operating systems, when FIPS support is enabled.

Problem:

- When attempting to display the Yarn Queue Manager interface, Cloudera Manager displays an error: "HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA".

1. Edit the file [/etc/default/cloudera-scm-server]
2. Around line 28, modify the line that starts with

```
#export CMF_OVERRIDE_TLS_CIPHERS=...
```

3. Remove the comment mark #.
4. Remove all ciphers with "3DES" in the name.
5. Save the file.
6. Restart the Cloudera Manager Server service.

COMPX-4644: Queue capacity rounding problem when configuration is initially set via YARN

When setting the capacity scheduler configuration through the YARN/Cloudera Manager configuration, there may be capacity values that use multiple decimal places. This results in rounding/floating point precision discrepancies in the UI when trying to validate that all sibling capacities equal 100%. The UI looks like all the numbers add up to 100, but the validation still displays an error and does not allow to save the capacities. It is also observed that the capacity is being calculated as, for example, 99.9999999991 in the backend.

- Create queues within the UI, or
- Ensure that capacities configured through the Capacity Scheduler safety valve do not have more than one decimal place.

20202 Database migration after enabling opt-in migration

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 CDP 7.1.9, forcing migration from H2 to PostgreSQL database.
- Upgrade to CDP 7.1.9 CHF2, moving back to H2 database.
- Upgrade to CDP 7.1.9 SP1 with valid PostgreSQL connection details in Queue Manager configurations.

To avoid any issues during the upgrade to version CDP 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.

CDPD-56559: MapReduce jobs can intermittently fail during a rolling upgrade.

During a rolling upgrade between CDP versions 7.1.8 and 7.1.9, MapReduce jobs may fail with message, RuntimeException: native snappy library not available. Although the native Snappy compression library is not loaded, a checkmark displays to indicate that the Snappy compression library is loading for NodeManagers that are pending upgrades. This causes the MapReduce jobs that are associated with the NodeManagers to fail. After the upgrade, the jobs work as expected. This issue only impacts rolling upgrades from before CDP 7.1.9 to a higher version.

None.

COMPX-12021 Queue Manager configurations on Scheduler Configuration page are not working

When setting the following properties on the YARN Queue Manager UI, the properties are set in the capacity-scheduler.xml file which does not have any effect on YARN. The properties need to be set in the yarn-site.xml file, which does not happen when you set them through YARN Queue Manager.

- "Maximum Application Priority" – "yarn.cluster.max-application-priority"
- "Enable Monitoring Policies" – "yarn.resourcemanager.scheduler.monitor.enable"
- "Monitoring Policies" – "yarn.resourcemanager.scheduler.monitor.policies"



Note: You can also set this property on the YARN configuration page in Cloudera Manager as "Capacity Scheduler Preemption".

- "Preemption: Observe Only" –
"yarn.resourcemanager.monitor.capacity.preemption.observe_only"
- "Preemption: Monitoring Interval (ms)" –
"yarn.resourcemanager.monitor.capacity.preemption.monitoring_interval"
- "Preemption: Maximum Wait Before Kill (ms)" –
"yarn.resourcemanager.monitor.capacity.preemption.max_wait_before_kill"
- "Preemption: Total Resources Per Round" –
"yarn.resourcemanager.monitor.capacity.preemption.total_preemption_per_round"
- "Preemption: Over Capacity Tolerance" –
"yarn.resourcemanager.monitor.capacity.preemption.max_ignored_over_capacity"
- "Preemption: Maximum Termination Factor" –
"yarn.resourcemanager.monitor.capacity.preemption.natural_termination_factor"
- "Enable Intra Queue Preemption" – "yarn.resourcemanager.monitor.capacity.preemption.intra-queue-preemption.enabled"

1. In Cloudera Manager, select the YARN service.
2. Click the Configuration tab.
3. Search for yarn-site.xml.
4. Under YARN Service Advanced Configuration Snippet (Safety Valve) for yarn-site.xml, add the corresponding parameter and value you need.
5. Click Save Changes.
6. Restart the YARN services.

COMPX-6214: When there are more than 600 queues in a cluster, potential timeouts occur due to performance reasons that are visible in the Configuration Service.

The Cloudera Manager proxy timeout configuration field is added now. This issue is tracked in OPSAPS-60554. For this release, the timeout is increased from 20 seconds to 5 minutes. However, if this problem occurs, Cloudera recommends you to increase the proxy timeout value.

COMPX-5817: YARN Queue Manager UI is not able to present a view of pre-upgrade queue structure. Cloudera Manager Store is not supported and therefore YARN does not have any of the pre-upgrade queue structure preserved.

When a Data Hub cluster is deleted, all saved configurations are also deleted. All YARN configurations are saved in Cloudera Manager Store and this is yet to be supported in Data Hub and Cloudera Manager. Hence, the YARN queue structure is also lost when a Data Hub cluster is deleted or upgraded or restored.

COMPX-6628: Unable to delete single leaf queue assigned to a partition.

In the current implementation, you cannot delete a single leaf queue assigned to a partition.

For each non-default partition, perform the following for the single child leaf queue and its parent queues:

1. In Cloudera Manager, click Cluster YARN .

2. Click the **Configuration** tab.
3. Search for ResourceManager. In the Filters pane, under Scope, select ResourceManager.
4. Add the following in Capacity Scheduler Configuration Advanced Configuration Snippet (Safety Valve):

```
Name: yarn.scheduler.capacity.<queuePath>.accessible-node-labels.<partition>.capacity
Value: 0
```

Set the value to 0 in Percentage mode, and 0w in Weight mode, and [memory=0,vcores=0] in Absolute mode.

```
Name: yarn.scheduler.capacity.<queuePath>.accessible-node-labels.<partition>.maximum-capacity
Value: 100
```

Set the value to 100 in Percentage and Weight mode and [memory=0,vcores=0] in Absolute mode.

5. Adjust the capacities of the siblings of the parent queue for the same partition.
6. Click Save Changes.
7. Restart the active **ResourceManager** service for the changes to apply.
8. In **Cloudera Manager**, click Cluster YARN Queue Manager UI .
9. Delete the desired single child leaf queue.

COMPX-5240: Restarting parent queue does not restart child queues in weight mode

When a dynamic auto child creation enabled parent queue is stopped in weight mode, its static and dynamically created child queues are also stopped. However, when the dynamic auto child creation enabled parent queue is restarted, its child queues remain stopped. In addition, the dynamically created child queues cannot be restarted manually through the YARN Queue Manager UI either.

Delete the dynamic auto child creation enabled parent queue. This action also deletes all its child queues, both static and dynamically created child queues, including the stopped dynamic queues. Then recreate the parent queue, enable the dynamic auto child creation feature for it and add the required static child queues.

COMPX-5589: Unable to add new queue to leaf queue with partition capacity in Weight/Absolute mode

Scenario

1. You create one or more partitions.
2. Assign a partition to a parent with children
3. Switch to the partition to distribute the capacities
4. Create a new child queue under one of the leaf queues but the following error is displayed:

```
Error :
2021-03-05 17:21:26,734 ERROR
com.cloudera.cpx.server.api.repositories.SchedulerRepository:
Validation failed for Add queue
operation. Error message: CapacityScheduler configuration
validation failed:java.io.IOException:
Failed to re-init queues : Parent queue 'root.test2' have
children queue used mixed of weight
mode, percentage and absolute mode, it is not allowed, please
double check, details:
{Queue=root.test2.test2childNew, label= uses weight mode}.
{Queue=root.test2.test2childNew,
label=partition uses percentage mode}
```

To create new queues under leaf queues successfully, perform the following:

1. Switch to Relative mode
2. Create the required queues
3. Create the required partitions
4. Assign partitions and set capacities
5. Switch back to Weight mode

1. Create the entire queue structure
2. Create the required partitions
3. Assign partition to queues
4. Set partition capacities

COMPX-5264: Unable to switch to Weight mode on creating a managed parent queue in Relative mode

In the current implementation, if there is an existing managed queue in Relative mode, then conversion to Weight mode is not allowed.

To proceed with the conversion from Relative mode to Weight mode, there should not be any managed queues. You must first delete the managed queues before conversion. In Weight mode, a parent queue can be converted into managed parent queue.

COMPX-5549: Queue Manager UI sets maximum-capacity to null when you switch mode with multiple partitions

If you associate a partition with one or more queues and then switch the allocation mode before assigning capacities to the queues, an Operation Failed error is displayed because `max-capacity` is set to null.

After you associate a partition with one or more queues, in the YARN Queue Manager UI, click Overview [****PARTITION NAME****] from the drop-down list and distribute capacity to the queues before switching allocation mode or creating placement rules.

COMPX-4992: Unable to switch to absolute mode after deleting a partition using YARN Queue Manager

If you delete a partition (node label) which has been associated with queues and those queues have capacities configured for that partition (node label), the CS.xml still contains the partition (node label) information. Hence, you cannot switch to absolute mode after deleting the partition (node label).

It is recommended not to delete a partition (node label) which has been associated with queues and those queues have capacities configured for that partition (node label).

COMPX-1445: YARN Queue Manager operations are failing when Queue Manager is installed separately from YARN

If Queue Manager is not selected during YARN installation, Queue Manager operations are failing. Queue Manager says 0 queues are configured and several failures are present. That is because ZooKeeper configuration store is not enabled.

1. In Cloudera Manager, select the YARN service.
2. Click the Configuration tab.
3. Find the Queue Manager Service property.
4. Select the Queue Manager service that the YARN service instance depends on.
5. Click Save Changes.
6. Restart all services that are marked stale in Cloudera Manager.

COMPX-3329: Autorestart is not enabled for Queue Manager in Data Hub

In a Data Hub cluster, Queue Manager is installed with autorestart disabled. Hence, if Queue Manager goes down, It does not restart automatically.

If Queue Manager goes down in a Data Hub cluster, you must go to the Cloudera Manager Dashboard and restart the Queue Manager service.

Third-party applications do not launch if MapReduce framework path is not included in the client configuration

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

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.

CDH-49165: History link in ResourceManager web UI broken for killed Spark applications

When a Spark application is killed, the history link in the ResourceManager web UI does not work.

To view the history for a killed Spark application, see the Spark HistoryServer web UI instead.

CDH-6808: Routable IP address required by ResourceManager

ResourceManager requires routable host:port addresses for `yarn.resourcemanager.scheduler.address`, and does not support using the wildcard `0.0.0.0` address.

Set the address, in the form `host:port`, either in the client-side configuration, or on the command line when you submit the job.

YARN cannot start if Kerberos principal name is changed

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.

Queue Manager does not open on using a custom user with a default Kerberos principal

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.

COMPX-8687: Missing access check for getAppAttempts

When the Job ACL feature is enabled using Cloudera Manager (YARN Configuration Enable JOB ACL property), the `mapreduce.cluster.acls.enabled` property is not generated to all configuration files, including the `yarn-site.xml` configuration file. As a result, the ResourceManager process uses the default value of this property. The default property of `mapreduce.cluster.acls.enabled` is `false`.

Enable the Job ACL feature using the Advanced Configuration Snippet:

1. In Cloudera Manager select the YARN service.
2. Click Configuration.
3. Find the YARN Service MapReduce Advanced Configuration Snippet (Safety Valve) property.
4. Click the plus icon and add the following:
 - Name: `mapreduce.cluster.acls.enabled`
 - Value: `true`

5. Click Save Changes.

Unsupported Features

The following YARN features are currently not supported in Cloudera Data Platform:

- Application Timeline Server v2 (ATSv2)
- Auxiliary Services
- Container Resizing
- Distributed or Centralized Allocation of Opportunistic Containers
- Distributed Scheduling
- Docker on YARN (DockerContainerExecutor) on Data Hub clusters
- Fair Scheduler
- GPU support for Docker
- Hadoop Pipes
- Moving jobs between queues
- Native Services
- Pluggable Scheduler Configuration
- Queue Priority Support
- Reservation REST APIs
- Resource Estimator Service
- Resource Profiles
- (non-ZooKeeper) ResourceManager State Store
- Rolling Log Aggregation
- Shared Cache
- YARN Federation

Known Issues in Apache ZooKeeper

Learn about the known issues in Zookeeper, the impact or changes to the functionality, and the workaround.

Zookeeper-client does not use ZooKeeper TLS/SSL automatically

The command-line tool ‘zookeeper-client’ is installed to all Cloudera Nodes and it can be used to start the default Java command line ZooKeeper client. However even when ZooKeeper TLS/SSL is enabled, the zookeeper-client command connects to localhost:2181, without using TLS/SSL.

Workaround:

Manually configure the 2182 port, when zookeeper-client connects to a ZooKeeper cluster. The following is an example of connecting to a specific three-node ZooKeeper cluster using TLS/SSL:

```
CLIENT_JVMFLAGS="-Dzookeeper.clientCnxnSocket=org.apache.zoo
keeper.ClientCnxnSocketNetty -Dzookeeper.ssl.keyStore.locati
on=<path to your configured keystore> -Dzookeeper.ssl.keyStor
e.password=<the password you configured for the keystore> -
Dzookeeper.ssl.trustStore.location=<path to your configured
truststore> -Dzookeeper.ssl.trustStore.password=<the password
you configured for the truststore> -Dzookeeper.client.secu
re=true" zookeeper-client -server <your.zookeeper.server-1>:218
2,<your.zookeeper.server-2>:2182,<your.zookeeper.server-3>:2182
```

Behavioral Changes In Cloudera Runtime 7.3.1

Behavioral changes denote a marked change in behavior from the previously released version to this version of Cloudera Runtime.

Behavioral Changes in Atlas

Behavioral changes denote a marked change in behavior from the previously released version to this version of Apache Atlas.

Summary:

The Exclude SubTypes and Exclude Sub-classifications filters were removed from the **Table** tab of entity details.

Previous behavior:

Previously, the Exclude SubTypes and Exclude Sub-classifications filters were available from the **Table** tab in entity details. There were no properties being passed to these filters when you visited the entity details of the page.

New behavior:

The two unused filter checkboxes Exclude SubTypes and Exclude Sub-classifications from the **Table** tab of entity detail page were removed.

Summary:

Special character validation was added to glossary, term and category names in Apache Atlas.

Previous behavior:

The special characters ('@', '!', '<', '>') could be used in glossary, term and category name fields.

New behavior:

The special characters ('@', '!', '<', '>') are no longer accepted in glossary, term and category name fields by the validation introduced. Avoid using these characters when creating glossary names, glossary terms and category names.

Behavioral Changes in Hive

Behavioral changes denote a marked change in behavior from the previously released version to this version of Apache Hive.

Summary:

Timestamp Conversion Handling

Previous behavior:

In Cloudera Private Cloud Base 7.1.9 version, hive.strict.timestamp.conversion was set to true, enforcing strict failures when casting between timestamp and numeric types.

New behavior:

Starting from Cloudera Private Cloud Base 7.3.1 version, reverted the strict mode for timestamp to numeric conversions. CAST operations between timestamp and numeric types no longer fail by default.

Summary:

Change in the way compaction initiator and cleaner threads are handled

Previous behavior:

The compaction initiator and cleaner threads are enabled and disabled by setting the `hive.compact.or.initiator.on` property to 'true' or 'false'.

New behavior:

A new property `hive.compactor.cleaner.on` is introduced that allows you to selectively enable or disable the cleaner thread.

This property is not listed and is set to 'true' by default. Add the property to Hive Metastore Server Advanced Configuration Snippet (Safety Valve) for `hive-site.xml` in Cloudera Manager to have the same out-of-the-box experience as in the previous version.

Also, ensure that you set the property to 'true' for the compactor to run on the HMS instance.

Behavioral Changes in Impala

Behavioral changes denote a marked change in behavior from the previously released version to this version of Apache Impala.

Summary:

Skips scheduling runtime filters for PK-FK joins when the build scan has no predicate filter, the join is a full table scan, and the bloom filter has a high false positive probability.

Previous behavior:

Runtime filters were scheduled for all PK-FK joins, regardless of effectiveness.

New behavior:

Filters are skipped for PK-FK joins when the build scan is a full table scan without filters, and the bloom filter has a high false positive probability, improving performance. For more details see, [Skip Scheduling Bloom Filter](#)

Apache Jira: [IMPALA-12357](#)

Summary:

Skips LZ4 compression when sending row batches within the same process to improve efficiency.

Previous behavior:

Row batches were serialized, compressed, sent through KRPC, and then decompressed, even when the sender and receiver were in the same process.

New behavior:

LZ4 compression is skipped for row batches sent within the same process, reducing unnecessary work and improving performance.

Apache Jira: [IMPALA-12430](#)

Behavioral Changes in Knox

Behavioral changes denote a marked change in behavior from the previously released version to this version of Apache Knox.

Knox token impersonation config

Summary

Knox token service has been changed to use the identity assertion provider configuration for impersonation.

Previous behaviour

The token service had its own impersonation configuration.

New behaviour

The token service relies on the identity assertion provider for impersonation configuration.

PEM file name change

Summary

The name of the pem file generated through `knoxcli.sh` has been changed.

Previous behaviour

The name of the file was `gateway-identity.pem`.

New behaviour

The name of the file is now `gateway-client-trust.pem`.

Composite authorization provider misconfiguration

Summary

Composite authorization provider misconfiguration behavior

Previous behaviour

1. If `composite.provider.names` is empty, the topology would fail deployment.
2. If `composite.provider.names` has an invalid value, the topology would fail deployment.

New behaviour

1. Deployment succeeds, and Knox allows access with no authorization since none is configured.
2. Deployment succeeds, but Knox rejects requests with a HTTP 403 response because the configuration is present (indicating that authorization is expected) but invalid.

Inactive topologies

Summary

Knox distinguishes inactive topologies from undeployed topologies.

Previous behaviour

Requests for topologies which are not yet fully deployed result in HTTP 404 responses.

New behaviour

Requests for topologies which are not yet fully deployed result in HTTP 503 responses.

Behavioral Changes in Livy

Behavioral changes denote a marked change in behavior from the previously released version to this version of Apache Livy.

Summary:

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

Previous behavior:

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

New behavior:

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

Behavioral Changes in Ranger

Behavioral changes denote a marked change in behavior from the previously released version to this version of Apache Ranger.

Summary: Ranger access audit behavior changes.

Previous behavior:

When you ran `hdfs dfs -copyFromLocal` command, audit logs were generated for the following:

- "write" Access Type and "write" permission.
- "rename" Access Type and "write" permission.
- "rename" Access Type and "write" permission.

When you ran `hdfs dfs -touch` command, audit log was generated for the following:

- "write" Access Type and "write" permission.

New behavior:

When you run `hdfs dfs -copyFromLocal` command, audit logs are generated for the following:

- "create" Access Type and "write" permission.
- "rename" Access Type and "write" permission.

When you run `hdfs dfs -touch` command, audit log is generated for the following:

- "create" Access Type and "write" permission.

Summary: Storagehandler authorisation has to be enabled for Ranger by setting the property "hive.security.authorization.tables.on.storagehandlers" to True in hive-site.xml file in HiveServer2 service.

Previous behavior:

This property was set to true by default.

New behavior:

In Data Hub, you configure `hive.security.authorization.tables.on.storagehandlers = true` to enable authorization of StorageHandler-based tables:

1. In Cloudera Manager, click **Clusters** **Hive** **Configurations**, and search for `hive.security.authorization.tables.on.storagehandlers`.
2. Set the value to true.
3. Save changes.

Behavioral Changes in Spark

Behavioral changes denote a marked change in behavior from the previously released version to this version of Apache Spark.

Summary:

Spark 2 has been removed from Cloudera Runtime.

Previous behavior:

Spark 2 was the default version in Cloudera Runtime, Spark 3 was available as an add-on parcel.

New behavior:

Spark 3 is the default Spark version in Cloudera Runtime. Spark 2 has been removed and no longer available in 7.3.1.0.



Important:

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

Refer to *Upgrading Spark* for more information on upgrading Spark clusters to 7.3.1.0, and *Migrating Spark Applications* for more information on migrating your existing Spark applications between versions 2 and 3.

Related Information

[Upgrading Spark](#)

[Migrating Spark Applications](#)

Deprecation Notices In Cloudera Runtime 7.3.1

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 CDP release. Marking an item as deprecated gives you time to plan for removal in a future CDP release.

Moving

Technology that Cloudera is moving from a future CDP 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 CDP release and plan for the alternative Cloudera offering or subscription for the technology.

Removed

Technology that Cloudera has removed from CDP 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.

Related Information

[Upgrading Spark](#)

[Migrating Spark Applications](#)

Platform and OS

The listed Operating Systems, databases, and instant client library are deprecated or removed from the 7.3.1 release.

Database Support:

The listed databases are deprecated from the 7.3.1 release:

- None

The following database is removed and no longer supported from the 7.3.1 release:

- None

Operating System

The listed operating systems are deprecated from the 7.3.1 release:

- None

The following operating system is removed and no longer supported from the 7.3.1 release:

- CentOS



Note: CentOS Linux 7 has reached end of life. Ensure to migrate to RHEL/Oracle Linux or any supported operating system before upgrading to 7.3.1.

Deprecation Notices for Apache Kafka

Certain features and functionality in Apache Kafka are deprecated or removed in Cloudera Runtime 7.3.1. 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 CDP. For a full list of deprecation and/or removals in the version Apache Kafka shipped with Runtime, review the *Notable Changes* as well as the *Release Notes* on <https://kafka.apache.org/>.

Deprecated

MirrorMaker (MM1)

MirrorMaker is deprecated. Cloudera recommends that you use Streams Replication Manager (SRM) 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 Apache Livy

Certain features and functionality in Apache Livy are deprecated or removed in Cloudera Runtime 7.3.1. You must review these changes along with the information about the features in Livy that will be removed or deprecated in a future release.

Removed

Apache Livy 2

As Spark 3 is the default Spark version in Cloudera Runtime, Livy 2 has been removed, alongside with Spark 2, and no longer available in 7.3.1

**Important:**

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

For more information on upgrading to Spark 3, refer to [Upgrading Spark](#) for more information on upgrading Spark clusters to 7.3.1, and [Migrating Spark Applications](#) for more information on migrating your existing Spark applications between versions 2 and 3.

Navigator Key Trustee Server

Certain features and functionality in Navigator Key Trustee Server (KTS) are deprecated or removed in CDP Private Cloud Base 7.3.1. You must review these changes along with the information about the features that will be removed or deprecated in a future release.

Navigator Key Trustee Server has been completely deprecated from UCL 7.3.1 release. For identical key and encryption management, customers must move to Ranger KMS. Also, customers using Navigator Encrypt (NavEncrypt) need to migrate metadata storage to Ranger KMS.

Deprecation Notices for Apache Oozie

Certain features and functionality in Apache Oozie are deprecated or removed in Cloudera Runtime 7.3.1. You must review these changes along with the information about the features in Oozie that will be removed or deprecated in a future release.

Deprecated

Oozie's Spark action

Due to the discontinuation and deprecation of Spark 2 in CDP 7.3.1, Cloudera decided to deprecate Oozie Spark actions, which are based on Spark 2. Consequently, Oozie's Spark actions will no longer be available, and if you attempt to execute a Spark action, an error will be raised.

Starting from 7.3.1, you must migrate to Spark 3 to use Spark actions. For more information, see [Spark 3 support in Oozie](#).

Deprecation Notices for Apache Spark

Certain features and functionality in Apache Spark 2 are deprecated or removed in Cloudera Runtime 7.3.1. You must review these changes along with the information about the features in Spark 2 that will be removed or deprecated in a future release.

Removed

Apache Spark 2

Spark 3 is the default Spark version in Cloudera Runtime. Spark 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, and [Migrating Spark Applications](#) for more information on migrating your existing Spark applications between versions 2 and 3.

Deprecation Notices for Apache Zeppelin

Certain features and functionality in Apache Zeppelin are deprecated or removed in Cloudera Runtime 7.3.1. You must review these changes along with the information about the features in Zeppelin that will be removed or deprecated in a future release.

Removed

Apache Zeppelin

Apache Zeppelin is removed from Cloudera Private Cloud.

You can reinstall the Zeppelin service as an external CSD and restore and use old notebooks, but Cloudera does not provide for Zeppelin starting in 7.3.1 and above. For more information, see [Reinstall Apache Zeppelin in 7.3.1](#)

Cloudera recommends you back up all existing Zeppelin notebooks before upgrading to version 7.3.1.

Fixed Common Vulnerabilities and Exposures 7.3.1

Common vulnerabilities and Exposures (CVEs) fixed in this release.

- [CVE-2023-6378](#) Logback
- [CVE-2023-6481](#) Logback
- [CVE-2023-2976](#) Google Guava
- [CVE-2020-8908](#) Google Guava
- [CVE-2018-10237](#) Google Guava
- [CVE-2023-52428](#) Nimbus-jose-jwt
- [CVE-2023-45865](#) Akka-actor
- [CVE-2021-42697](#) Akka-http-core
- [CVE-2021-23339](#) Akka-http-core
- [CVE-2022-31023](#) Akka-http-server
- [CVE-2021-26291](#) Apache Maven
- [CVE-2022-46337](#) Apache Derby
- [CVE-2023-22006](#) Graal-sdk
- [CVE-2023-50386](#) Apache Solr
- [CVE-2023-50291](#) Apache Solr
- [CVE-2023-50292](#) Apache Solr
- [CVE-2023-50298](#) Apache Solr
- [CVE-2023-1932](#) Hibernate Validator
- [CVE-2024-22201](#) Eclipse Jetty
- [CVE-2024-21634](#) Amazon Ion
- [CVE-2017-7525](#) Jackson-mapper-asl
- [CVE-2019-10172](#) Jackson-mapper-asl
- [CVE-2023-51775](#) Jose4j
- [CVE-2020-15522](#) Bouncycastle
- [CVE-2020-0187](#) Bouncycastle
- [CVE-2022-1471](#) Snakeyaml
- [CVE-2022-25857](#) Snakeyaml
- [CVE-2022-38749](#) Snakeyaml
- [CVE-2022-38751](#) Snakeyaml

- [CVE-2022-38752](#) Snakeyaml
- [CVE-2022-41854](#) Snakeyaml
- [CVE-2022-38750](#) Snakeyaml
- [CVE-2021-31684](#) Json-smart
- [CVE-2023-1370](#) Json-smart
- [CVE-2021-27568](#) Json-smart
- [CVE-2021-4178](#) Fabric 8 Kubernetes client
- [CVE-2023-3635](#) Okio
- [CVE-2024-1597](#) Postgresql
- [CVE-2023-45857](#) Axios
- [CVE-2022-4244](#) Plexus-utils
- [CVE-2022-4245](#) Plexus-utils
- [CVE-2023-34453](#) Snappy-java
- [CVE-2023-34454](#) Snappy-java
- [CVE-2023-34455](#) Snappy-java
- [CVE-2023-43642](#) Snappy-java
- [CVE-2023-34042](#) Spring Security
- [CVE-2024-22257](#) Spring Security
- [CVE-2023-20859](#) Spring Vault
- [CVE-2024-22243](#) Spring Framework
- [CVE-2024-22262](#) Spring Framework
- [CVE-2024-22259](#) Spring Framework
- [CVE-2024-1300](#) Vertx-core
- [CVE-2023-44483](#) Xmlsec
- [CVE-2024-31573](#) Xmlunit-core
- [CVE-2024-38998](#) Requirejs
- [CVE-2024-38999](#) Requirejs
- [CVE-2023-4759](#) Eclipse Jgit