

Cloudera Runtime 7.3.1

Release Notes

Date published: 2024-12-11

Date modified: 2025-11-06

CLOUdera

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

Legal Notice

© Cloudera Inc. 2025. All rights reserved.

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

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

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

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

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

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

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

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

Contents

Overview.....	7
What's New In Cloudera Runtime 7.3.1.....	7
Atlas.....	7
What's New in Cloud Connectors.....	8
Cruise Control.....	9
HBase.....	9
HDFS.....	11
Hive.....	12
Impala.....	15
Hue.....	19
Iceberg.....	21
Kafka.....	23
Kerberos.....	25
Knox.....	26
Kudu.....	26
Livy.....	27
Navigator Encrypt.....	27
Oozie.....	28
Ozone.....	29
Phoenix.....	29
Ranger.....	30
What's New in Ranger KMS.....	32
Schema Registry.....	32
Solr.....	33
Spark.....	34
Sqoop.....	35
Streams Messaging Manager.....	35
Streams Replication Manager.....	37
YARN and YARN Queue Manager.....	39
ZooKeeper.....	39
What's new in Platform Support.....	40
Cloudera Runtime Component Versions.....	41
Cloudera Runtime 7.3.1.500 SP3.....	41
Cloudera Runtime 7.3.1.400 SP2.....	42
Cloudera Runtime 7.3.1.300 SP1 CHF 1.....	44
Cloudera Runtime 7.3.1.200 SP1.....	45
Cloudera Runtime 7.3.1.100 CHF 1.....	47
Cloudera Runtime 7.3.1.....	48
Using the Cloudera Runtime Maven repository 7.3.1.....	50
Cloudera Runtime 7.3.1.500-182.....	50
Cloudera Runtime 7.3.1.400-100.....	68

Cloudera Runtime 7.3.1.300-81.....	86
Cloudera Runtime 7.3.1.200-90.....	104
Cloudera Runtime 7.3.1.100-57.....	122
Cloudera Runtime 7.3.1.0-197.....	140

Release Matrix.....158

Fixed Issues In Cloudera Runtime 7.3.1..... 159

Atlas.....	159
Avro.....	163
Cloud Connectors.....	163
Cruise Control.....	164
HDFS.....	165
HBase.....	166
Hive.....	172
Impala.....	187
Hue.....	196
Fixed Issues in Apache Iceberg.....	200
Kafka.....	204
Kudu.....	205
Knox.....	207
Livy.....	210
Navigator Encrypt.....	211
Oozie.....	212
Ozone.....	213
Parquet.....	221
Phoenix.....	222
Ranger.....	223
Ranger KMS.....	230
Schema Registry.....	231
Solr.....	232
Spark.....	234
Spark Atlas Connector.....	236
Sqoop.....	236
Streams Messaging Manager.....	237
Streams Replication Manager.....	238
Yarn and Yarn Queue Manager.....	239
ZooKeeper.....	241

Known Issues In Cloudera Runtime 7.3.1..... 241

Known Issues in Platform.....	241
Atlas.....	242
Avro.....	250
Known Issues in Cloud Connectors.....	251
Cruise Control.....	252
HBase.....	253
HDFS.....	255
Hive.....	258
Impala.....	259
Hue.....	264
Iceberg.....	268
Kafka.....	271
Knox.....	275

Kudu.....	278
Livy.....	279
Navigator Encrypt.....	279
Oozie.....	280
Ozone.....	281
Parquet.....	290
Phoenix.....	291
Ranger.....	291
Ranger KMS.....	294
Schema Registry.....	295
Solr.....	296
Spark.....	305
Spark Atlas Connector.....	306
Sqoop.....	307
Streams Messaging Manager.....	308
Streams Replication Manager.....	310
Yarn and Yarn Queue Manager.....	312
ZooKeeper.....	319

Behavioral Changes In Cloudera Runtime 7.3.1..... 319

Behavioral Changes in Atlas.....	320
Behavioral Changes in Apache Avro.....	321
Behavioral Changes in Cloud Connectors.....	322
Behavioral Changes in Cruise Control.....	322
Behavioral Changes in HBase.....	323
Behavioral Changes in HDFS.....	323
Behavioral Changes in Hive.....	324
Behavioral Changes in Impala.....	327
Behavioral Changes in Hue.....	329
Behavioral Changes in Iceberg.....	330
Behavioral Changes in Kafka.....	330
Behavioral Changes in Knox.....	331
Behavioral Changes in Kudu.....	332
Behavioral Changes in Livy.....	333
Behavioral Changes in Oozie.....	334
Behavioral Changes in Ozone.....	334
Behavioral Changes in Apache Parquet.....	335
Behavioral Changes in Phoenix.....	335
Behavioral Changes in Ranger.....	336
Behavioral Changes in Schema Registry.....	337
Behavioral Changes in Solr.....	338
Behavioral Changes in Sqoop.....	338
Behavioral Changes in Streams Messaging Manager.....	339
Behavioral Changes in Streams Replication Manager.....	339
Behavioral Changes in Spark.....	340
Behavioral Changes in Spark Atlas Connector.....	342
Behavioral Changes in Yarn and Yarn Queue Manager.....	342

Deprecation Notices In Cloudera Runtime 7.3.1..... 343

Platform and OS.....	344
Deprecation Notices for Apache Kafka.....	345
Deprecation Notices for Apache Livy.....	345
Navigator Key Trustee Server.....	345
Deprecation Notices for Apache Oozie.....	346

Deprecation Notices for Apache Spark.....	346
Deprecation Notices for Apache Zeppelin.....	346
Technical Service Bulletins.....	347
TSB 2025-820: Potential Data Integrity Issues Found in Ozone.....	347
TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager.....	347
Apache Parquet CVE-2025-30065.....	348
Fixed Common Vulnerabilities and Exposures.....	349
Cloudera Runtime 7.3.1.500 SP3.....	349
Cloudera Runtime 7.3.1.400 SP2.....	349
Cloudera Runtime 7.3.1.300 SP1 CHF 1.....	350
Cloudera Runtime 7.3.1.200 SP1.....	350
Cloudera Runtime 7.3.1.100 CHF 1.....	350
Cloudera Runtime 7.3.1.....	350

Overview

This document provides a summary of the latest updates in Cloudera Runtime 7.3.1 and its Service Packs and Cumulative Hotfixes. It includes new features, improvements, known and fixed issues, technical previews, and more. For detailed, component-level information, see the [Cloudera documentation](#).

Understanding Cloudera platform versioning

Starting with Cloudera 7.3.1.0, versioning follows the <Base>.<Major>.<Minor>.<Patch> format:

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

For example, in version **7.3.1.0**, **7** is the base version, **3** is the major version, **1** is the minor version, and **0** indicates that it is an initial minor release (not a CHF or SP). For more details, see the [Cloudera Platform Support Policy](#).

What's New In 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.



Note: For information about Cloudera Base on premises 7.3.1 service packs upgrade paths, see [Supported in-place upgrade paths](#)

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, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

Improved Atlas metrics API performance

The performance of Atlas API service metrics is significantly enhanced by making detailed entity type counts optional. For large Atlas deployments with many entity types, API response times are now much faster by default, while still allowing users to request the complete metrics when needed for detailed analysis.

Key changes:

- Enhanced the `MetricsService.getMetrics` method to accept an `excludeTypeAndSubTypeEntity` boolean parameter, allowing callers to exclude metrics for entity types and their subtypes.
- Updated logic to conditionally calculate and include metrics for entity types and subtypes only when requested.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1:

There are no new features in this release.

What's New in Cloud Connectors

Learn about the new features of Cloud Connectors Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

There are no new features this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features this release.

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 AWS Region to store your data. This enables you to have your storage and compute resources 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.

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

Vectorized IO support

Support for Hadoop Vectors 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.

What's New in Cruise Control

Learn about the new features of Cruise Control in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

There are no new features in this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

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 Streams Messaging Manager using the  on the navigation sidebar.

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

What's New in Apache HBase

Learn about the new features of Apache HBase in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

There are no new features in this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1:**HBase supports JDK 17**

HBase supports Oracle JDK version 17.0.6 starting from CDP Runtime 7.1.9. For more information on JDK 17, see [Java Requirements](#).

HBase rebase to 2.4.17

CDP Private Cloud Base is updated to use Apache HBase version 2.4.17 and Apache HBase Thirdparty to base version 4.1.1 for a smoother and better functionality. Upgrade your HBase client applications for seamless connectivity.

HBase supports Snappy with /tmp directory mounted with noexec option

In Cloudera Manager, the Snappy temporary directory configuration item is added to HBase Master and HBase RegionServer to allow Snappy compression when /tmp directory is mounted with noexec option.

Operating system support

HBase is now supported on the following operating systems:

- RHEL-9.1
- RHEL-8.8
- RHEL-8.8 FIPS
- Oracle-8.8 UEK
- SLES-15 SP4 for x86

HBase supports load balancing using a cache-aware load balancer

The HBase balancer now supports the cache-aware load balancer that enhances the capability of HBase to enable the balancer to consider the cache allocation of each region on region servers while calculating a new assignment plan. This balancer also uses the region or region server cache allocation information reported by the region servers to calculate the percentage of HFiles cached for each region on the hosting server, and then use that as an additional factor while deciding an optimal new assignment plan.

HBase supports Snappy with /tmp directory mounted with noexec option

In Cloudera Manager, the Snappy temporary directory configuration item is added to HBase Master and HBase RegionServer to allow Snappy compression when /tmp directory is mounted with noexec option.

HBase supports Netty native libraries with /tmp directory mounted with noexec option

In Cloudera Manager, the Netty native library working directory configuration item is added to HBase Master and HBase RegionServer to support HBase with /tmp directory mounted with noexec option.

HBase shows cached percentage for region data on RegionServer UI

An important feature for Cloudera Operational Database (COD) over S3 with ephemeral cache is the process of warming up the cache at region opening (also known as *cache prefetch*). The goal is to load the most of the dataset before any client reads, so that a reduced latency and optimal performance can be achieved for the application requests. This *prefetch* process takes several hours on very large datasets, and the operators might want to monitor the progress of this cache loading. To handle this, HBase has introduced new metrics about the percentage of individual regions data currently cached, and it also added this information to the Storefile Metrics tab in the Regions section of the RegionServer UI.

Related Apache JIRA: [HBASE-28246](#)

Region Name	Num. Stores	Num. Storefiles	Storefile Size Uncompressed	Storefile Size	Index Size	Bloom Size	Data Locality	% Cached
hbase:namespace,,1710343972022.123e5d1573867d44a01e1460a5dd6e8.	1	1	1 MB	1 MB	1 KB	1 KB	1.0	100.00%
my-user-table.A1009037-128-dup1709753704318.1709757279182.64e3c9c785002d1c88ba1f6c5e7527.	2	2	201 MB	32 MB	31 KB	256 KB	1.0	100.00%
my-user-table.A1009093-dup1709748054721.1709759852911.a34c576f6e6ad1f778e04e7c343579d.	2	2	405 MB	66 MB	63 KB	512 KB	1.0	100.00%
my-user-table.A1009238-dup1709753529701.1709758902574.cd99cd790caa93522d9beb313d14b0c0.	2	2	402 MB	65 MB	63 KB	512 KB	1.0	100.00%
my-user-table.A1009226-208-dup1709749478903.1709756783373.9e1abc072025ca6b0fc7c945212525a1.	2	2	399 MB	65 MB	62 KB	512 KB	1.0	98.46%

HBase supports disabling the caching for the individual column families

In some use cases, not all tables in the dataset have the same SLA requirements. If the total cache capacity is much smaller than the whole dataset, an alternative is to restrict the cache usage by the tables with critical response times. In HBase, you can now implement this by disabling the cache on individual column families.

On an hbase shell, perform the following `alter` command for each column family that does not require caching.

```
alter 'NAMESPACE:TABLENAME', {NAME=>'CF_NAME', BLOCKCACHE => 'false'}
```

HBase supports truncating the regions in a table

You can now truncate individual regions of an HBase table using the `truncate_region` command.

The command syntax is as follows.

```
truncate_region 'REGIONNAME'
```

```
truncate_region 'ENCODED_REGIONNAME'
```

For example,

```
hbase:008:0> list_regions 'employee'
                                SERVER_NAME
                                REGION_NAME |
START_KEY | END_KEY | SIZE | REQ | LOCALITY |
-----|-----|-----|-----|-----|
|-----|-----|-----|-----|-----|
ccycloud-4.nightly-7x-by.root.comops.site,22101,1718869191555 |
employee,,1718877308795.66828b0fe6ceda3e28608617eb6f6b3f. |
| 2 | 1 | 2 | 1.0 |
ccycloud-2.nightly-7x-by.root.comops.site,22101,1718869191308 |
employee,2,1718877308795.ff9b19452fecea6353694583e3473b5b. |
2 | 1 | 2 | 1.0 |
2 rows
Took 0.1088 seconds
hbase:014:0> truncate_region 'employee,2,1718877308795.ff9b19452fecea6353694583e3473b5b.'
Took 0.6236 seconds
hbase:010:0> truncate_region 'ff9b19452fecea6353694583e3473b5b'
Took 0.6500 seconds
```

What's New in Apache Hadoop HDFS

There are no new features for Apache Hadoop HDFS in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

There are no new features in this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1:

For more information about HDFS, see [HDFS Overview](#)

What's New in Apache Hive

Learn about the new features of Hive in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:**Common table expression detection and rewrites using cost-based optimizer**

Hive's existing shared work optimizer detects and optimizes common table expressions heuristically, but it lacks cost-based analysis and has limited customization. New APIs and configuration options have been introduced to support common table expression optimizations at the cost-based optimizer level. This feature is experimental and disabled by default.

Apache Jira: [HIVE-28259](#)

Upgraded Avro to version 1.11.3**OpenTelemetry integration for Hive**

Hive now integrates with OpenTelemetry (OTel) to enhance query by collecting and exporting telemetry data, including infrastructure and workload metrics. An OTel agent in Cloudera helps monitor query performance and troubleshoot failures.

Apache Jira: [HIVE-28504](#)

For more information, see [OpenTelemetry support for Hive](#)

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

Configurable TCP keepalive for client connections

New startup options allow you to enable and tune TCP keepalive settings for client connections, helping detect dead clients and prevent premature disconnections caused by network disruptions or load balancer timeouts. The new startup options include:

- `client_keepalive_probe_period_s`: Idle time before sending keepalive probes. If set to `> 0`, keepalive is enabled.
- `client_keepalive_retry_period_s`: Time between keepalive probes.
- `client_keepalive_retry_count`: Number of keepalive probes.

New features in Cloudera Runtime 7.3.1

Trusting HTTP headers for authentication

When HTTP headers are authenticated through 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 Apache Impala

Learn about the new features of Apache Impala in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

Python 3.9 Support for Impala on RHEL 8

Cloudera now provides support for Python 3.9 on RHEL 8 within Cloudera 7.3.1. This new capability ensures that Impala components dependent on Python libraries function completely and reliably.

SHOW VIEWS statement

This release introduces the [SHOW VIEWS statement](#), which simplifies the task of listing all views within a specified schema or database. Using this command, you can quickly identify and review views, thereby enhancing performance by reducing metadata scan operations.

Planner changes to improve cardinality estimation

Significant changes have been made to the [query planner to improve cardinality estimation](#) , a critical component of workload-aware autoscaling.

In previous versions, Impala would generate a plan first and then search for runtime filters based on the entire plan. In this release, selective runtime filters have been integrated. These filters aim to reduce the cardinality estimates of scan nodes and specific join nodes located above them. This refinement occurs after the generation of runtime filters and before the computation of resource requirements.

Distribute runtime filter aggregation

[Aggregating runtime filters](#) during runtime can impose significant memory overhead on the coordinator. To address this issue, we have implemented a solution that distributes the runtime filter aggregation across specific Impala backends.

Improvement in catalog observability

This release introduces significant enhancements to the [Impala Catalog Web UI](#) , focusing on addressing performance issues associated with delays in processing Hive Metastore (HMS) events. These improvements aim to mitigate the risk of queries using outdated metadata.

Caching codegen functions

In Impala, "codegen" involves generating specialized machine code for each query based on query-specific information. The [codegen capability](#) converts query-specific information into machine code, enhancing query performance through faster execution.

Support ORDER BY for collections of variable length types in SELECT list

This release introduces support for collections of variable length types in the sorting tuple. While it's now possible to include these collection columns in the SELECT list alongside other columns used for sorting, direct sorting by these collection columns is not supported. Additionally, collections of variable-length types can now serve as non-passthrough children of UNION ALL nodes.

It is important to note that structs containing collections, whether of variable or fixed length, are still not supported in the select list for ORDER BY queries.

Here are examples of supported queries:

```
select id, arr_string_id from collection_tbl order by id;  
select id, map_id from collection_tbl order by id;
```

However, queries such as the following are not supported:

```
select id, struct_contains_map from collection_struct_mix order  
by id;
```

Attempting to execute such queries will result in the error message `AnalysisException: Sorting is not supported if the select list contains collection(s) nested in struct(s)`.

Improved cardinality estimation for aggregation queries

Impala now provides more accurate cardinality estimates for aggregation queries by considering data distribution, predicates, and tuple tracing. Enhancements include:

- **Pre-aggregation Cardinality Adjustments:** A new estimation model accounts for duplicate keys across nodes, reducing underestimation errors.
- **Predicate-Aware Cardinality Calculation:** The planner now considers filtering conditions on group-by columns to refine cardinality estimates.
- **Tuple Tracing for Better Accuracy:** Improved tuple analysis allows deeper tracking across views and intermediate aggregation nodes.
- **Consistent Aggregation Node Stats Computation:** The planning process now ensures consistent and efficient recomputation of aggregation node statistics. These improvements lead to better memory estimates, optimized query execution, and more efficient resource utilization.

- Tuple-Based Cardinality Analysis: Analyzing grouping expressions from the same tuple to ensure their combined number of distinct values does not exceed the output cardinality of the source PlanNode, reducing overestimation.
- Refined number of distinct values Calculation for CPU Costing: The new approach applies a probabilistic formula to a single global NDV estimate, improving accuracy and reducing overestimation in processing cost calculations.

Apache Jira: [IMPALA-2945](#), [IMPALA-13086](#), [IMPALA-13465](#) , [IMPALA-13526](#), [IMPALA-13405](#), [IMPALA-13644](#)

Cleanup of host-level remote scratch dir on startup and exit

Impala now removes leftover scratch files from remote storage during startup and shutdown, ensuring efficient storage management. The cleanup targets files in the host-specific directory within the configured remote scratch location.

A new flag, `remote_scratch_cleanup_on_start_stop`, controls this behavior. By default, cleanup is enabled, but you can disable it if multiple Impala daemons on a host or multiple clusters share the same remote scratch directory to prevent unintended deletions.

Apache Jira: [IMPALA-13677](#), [IMPALA-13798](#)

Graceful shutdown with query cancellation

Impala now attempts to cancel running queries before reaching the graceful shutdown deadline, ensuring resources are released properly. The new `shutdown_query_cancel_period_s` flag controls this behavior. The default value is 60 seconds. If set to a value greater than 0, Impala will try to cancel running queries within this period before forcing shutdown. If the value exceeds 20% of the total shutdown deadline, it is automatically capped to prevent excessive delays. This approach helps prevent unfinished queries and unreleased resources during shutdown. For more information, see [Setting Impala Query Cancellation on Shut down](#)

Programmatic query termination

Impala now supports the KILL QUERY statement, enabling you to forcibly terminate queries for better workload management. The KILL QUERY statement cancels and unregisters queries on any coordinator. For more information, see [KILL QUERY statement](#)

AI Functions in Impala

Cloudera Runtime introduces Impala's built-in `ai_generate_text` function integrates Large Language Models (LLMs) into SQL for tasks such as sentiment analysis and translation. It simplifies workflows, requires no ML expertise, and supports default or custom UDF configurations.

Secure API key storage is supported through a JCEKS keystore. A lightweight tool included in the UDF SDK helps create or update keystores on Amazon S3 or Azure ABFS without a local Hadoop setup.

For more information, see [Advantages and use cases of Impala AI functions](#)

Ability to log and manage Impala workloads (Technical Preview)

Cloudera Runtime provides you the option to enable logging Impala queries on an existing Virtual Warehouse or while creating a new Impala Virtual Warehouse. The information for all completed Impala queries is stored in the `sys.impala_query_log` system table. Information about all actively running and recently completed Impala queries is stored in the `sys.impala_query_live` system table. Users with appropriate permissions can query this table using SQL to monitor and optimize the Impala engine.

For more information, see [Impala workload management](#)

Support for Impala external JDBC data sources (Technical Preview)

Apache Impala now supports reading from external JDBC data sources. An external JDBC table represents a table or a view in a remote RDBMS database or another Impala cluster. Using external JDBC tables, you can connect Impala to a database, such as MySQL, PostgreSQL, or another Impala cluster and read the data in the remote tables.

For more information, see [Using Impala to query external JDBC data sources](#)

Running queries on system tables (Technical Preview)

Queries against Impala system tables, such as `sys.impala_query_live`, could get delayed due to admission control constraints. These queries, which require only coordinator resources, were previously blocked by queries competing for executor resources. To address this, Impala introduces an "only coordinators" request pool, allowing system table queries to bypass executor queues and run only on the coordinators to prevent delays during admission.

Apache Jira: [IMPALA-13201](#)

For more information, see [Running queries on system tables](#)

User quotas in admission control (Technical Preview)

This release introduces user quotas in Impala admission control, a new feature designed to enhance resource management and ensure fair query distribution across users and groups.

For more information, see [User quotas in Admission Control](#)

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1

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: [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](#)

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 Hue

Learn about the new features of Hue in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:**Multi-Python Support for Hue**

This release introduces Python 3.11 support for Hue across various operating systems. This enhancement ensures continued cluster stability by offering concurrent support for multiple Python versions, including Python 3.8 and Python 3.9. To access Python 3.11 for Hue, you must upgrade to **Cloudera Manager 7.13.1.500** version.

You can now upgrade to newer Python versions without disrupting ongoing cluster operations, minimizing downtime, and ensuring a smooth transition.

Adopting Python 3.11 gives you access to adopt to the latest features, security patches, and performance improvements.

By default, Hue automatically selects the highest available Python version on the system. To set a specific lower version of Python, see [How to configure a specific Python version for Hue](#).

For more information, see [Python support matrix](#).

Hue SQL AI: Multi-database querying now supported

The Hue SQL AI Assistant now supports multi-database querying, allowing you to retrieve data from multiple databases simultaneously. This enhancement simplifies managing large datasets across different systems and enables seamless cross-database queries.

- Support for cross-database queries.
- Ability to retrieve and combine data from multiple sources in a single query.

For more information, see [Multi database support for SQL query](#).

User Input Validation for Hue SQL AI

Hue SQL AI now supports secure and optimized integration with large language models (LLMs). You can now configure user input validation, such as prompt length limits, regex restrictions, and HTML tag handling, and more, to enhance both security and system performance.

For more information, see [User Input Validation for Hue SQL AI](#).

New Hue Storage Browser (Technical Preview)

The Hue Storage Browser is a web-based interface designed to provide seamless interaction with multiple file systems. With enhanced usability and functionality, the Storage Browser improves data management within the Hadoop ecosystem, offering a streamlined experience.

For more information, see [Hue Storage Browser](#).

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:**Secure Approach for Passing a Token in Cloudera Manager**

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

For more information, see [Secure Approach for Passing a Token in Cloudera Manager](#).

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 Cloudera 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, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

Cloudera support for Apache Iceberg version 1.5.2

The Apache Iceberg component has been upgraded from 1.3.1 to 1.5.2.

Insert into partition and insert overwrite partition support

From Hive you can insert into, or overwrite data in, Iceberg tables that are statically or dynamically partitioned. For syntax and limitations, see [Insert into/overwrite partition support](#).

Truncate partition support

This release introduces the capability to [truncate an Iceberg partition](#) from Hive. Truncation removes all rows from the table and a new snapshot is created.

Enhancement of expiring Iceberg snapshots

The Expire Snapshots feature has been enhanced to offer more flexibility. In addition to expiring snapshots older than a timestamp, you can now expire snapshots based on the following conditions:

- A snapshot having a given ID
- Snapshots having IDs matching a given list of IDs
- Snapshots within the range of two timestamps

You can keep snapshots you are likely to need, for example recent snapshots, and expire old snapshots. For example, you can keep daily snapshots for the last 30 days, then weekly snapshots for the past year, then monthly snapshots for the last 10 years. You can remove specific snapshots to meet the GDPR right to be forgotten requirements.

For more information, see the [Expire snapshots feature](#).

Iceberg branching and tagging

From Hive, you can manage the lifecycle of snapshots using the [Iceberg branching](#) and [Iceberg tagging](#) features. Branches are references to snapshots that have a lifecycle of their own. Tags identify snapshots you need for auditing and conforming to GDPR.

Drop partition feature

You can easily [remove a partition](#) from an Iceberg partition using an alter table statement from Impala. Removing a partition does not affect the table schema. The column is not removed from the schema.

Support for copy-on-write (COW)

Hive supports the copy-on-write (COW) as well as merge-on-read (MOR) for handling Iceberg row-level updates and deletes. You [configure COW or MOR](#) based on your use case and rate of data change.

Support for Iceberg data compaction and related enhancements

You can compact Iceberg tables and optimize them for read operations from Hive and Impala. Compaction is an essential table maintenance activity that creates a new snapshot, which contains the table content in a compact form.

The OPTIMIZE TABLE statement also includes the following improvements:

- Supports partition evolution

The Hive and Impala OPTIMIZE TABLE supports compaction of Iceberg tables with partition evolution.

- Supports data compaction based on file size threshold

The Impala OPTIMIZE TABLE statement includes a FILE_SIZE_THRESHOLD_MB option that enables you to specify the maximum size of files (in MB) that should be considered for compaction.

For more information, see [Iceberg data compaction](#).

SQL support for querying Iceberg metadata tables

Apache Iceberg stores extensive metadata for its tables. From Hive and Impala, you can query the metadata tables as you would query a regular table. For example, you can use projections, joins, filters, and so on. See [Query metadata tables feature](#).

Directed distribution mode

This release implements directed distribution mode. The scheduler collects information about which Iceberg data file is scheduled on which host. Since, the scan node for the data files are on the same host as the Iceberg join node, delete files are sent directly to that specific host. This mode can improve V2 table read performance.

Impala support for reading Iceberg equality deletes for NiFi (Preview)

Cloudera supports row-level deletes, and starting with this release you can read equality deletes from Impala with support added for Apache NiFi. See the [Delete data feature](#).

Reading Iceberg Puffin statistics

Impala supports reading Puffin statistics from current and older snapshots. When there are Puffin statistics for multiple snapshots, Impala chooses the most recent statistics for each column. This indicates that statistics for different columns may come from different snapshots. If there are Hive Metastore (HMS) and Puffin statistics for a column, the most recent statistics are considered. For HMS statistics, the `impala.lastComputeStatsTime` property is used and for Puffin statistics, the snapshot timestamp is used to determine which among the two is the most recent. For more information, see [Iceberg Puffin statistics](#).



Note: Reading Puffin statistics is disabled by default. Set the `--enable_reading_puffin_stats` startup flag to "true" to enable it.

Impala supports the UPDATE statement for Iceberg tables

This release introduces support for the UPDATE statement in Impala for Iceberg tables. With this enhancement, you can now use Impala to update data in a V2 Iceberg table. The UPDATE operations in Impala are executed as an atomic pair of DELETE and INSERT operations. The Iceberg V2 format supports row-level modifications using delete files, enabling seamless row-level updates. For more information, see the [Iceberg Update data feature](#)

Impala supports the MERGE INTO statement for Iceberg tables

You can use Impala to run a MERGE INTO statement on an Iceberg table based on the results of a join between a target and source Iceberg table. For more information, see the [Iceberg Merge feature](#).

Cloudera Runtime 7.3.1.400 SP2

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF1

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF1

There are no new features in this release.

Cloudera Runtime 7.3.1**Apache Iceberg support for Hive**

Cloudera 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, Cloudera Base on premises 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 Cloudera, 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](#).

Cloudera 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 Kafka

Learn about the new features of Apache Kafka in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:**The LdapLoginModule is blocked by default**

The LdapLoginModule (com.sun.security.auth.module.LdapLoginModule) class is blocked by default and can not be used in SASL JAAS configurations. Use the PlainLoginModule (org.apache.kafka.common.security.plain.PlainLoginModule) class for LDAP authentication.

Configurable allow list of URLs for OAuth authentication

A new Java option, org.apache.kafka.sasl.oauthbearer.allowed.urls, is introduced for Kafka brokers and Kafka Connect that enable you to specify an allow list of URLs for the OAuth authentication mechanism. This gives you stricter control over which URLs can be accessed by Kafka for authentication.

By default the option is empty, which means that a connection to any URL is allowed. The configuration property you use in Cloudera Manager to set the option is different for Kafka brokers and Kafka Connect. Use the following configuration properties for Kafka brokers and Kafka Connect:

- For **Kafka brokers**, configure the option by adding it to the Additional Broker Java Options Kafka broker property in Cloudera Manager. For example:

```
-Dorg.apache.kafka.sasl.oauthbearer.allowed.urls=HTTP://  
WWW.OAUTH-EXAMPLE-1.COM,HTTP://WWW.OAUTH-EXAMPLE-2.COM
```

- For **Kafka Connect**, configure the option by adding it to the EXTRA_ARGS environment variable. The EXTRA_ARGS environment variable is configured by adding it to the Kafka Connect Environment Advanced Configuration Snippet (Safety Valve) property in Cloudera Manager. For example:

```
EXTRA_ARGS=-Dorg.apache.kafka.sasl.oauthbearer.allowed.u  
r ls="HTTP://WWW.OAUTH-EXAMPLE-1.COM,HTTP://WWW.OAUTH-  
EXAMPLE-2.COM"
```

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

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 Cloudera 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 Apache Knox

Learn about the new features of Apache Knox in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

There are no new features in this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

Improved Knox's auto-discovery feature for large clusters

Earlier, the auto-discovery feature of Cloudera Manager through Knox fetched role configuration parameters for each service role. This caused a large number of REST API calls in case of large clusters.

The user can now specify a new auto-discovery method in Knox, which improves the performance of Knox's auto-discovery feature for large clusters.

For more information, see [Configure Knox auto-discovery for large clusters](#).

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1

Improved custom descriptor and shared provider configuration management in Knox

The users can now remove previously created custom descriptors, including their associated topologies, and shared provider configurations when they are no longer needed.

For more information, see [Remove a custom descriptor from Apache Knox](#) and [Remove a shared provider configuration](#).

What's New in Apache Kudu

Learn about the new features of Apache Kudu in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

There are no new features in this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1:

There are no new features in this release.

What's New in Livy

Learn about the new features of Apache Livy in Cloudera Runtime 7.3.1 release, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in Livy in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in Livy in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in Livy in this release.

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 Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

There are no new features in this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

Cloudera Runtime 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, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

There are no new features in this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

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 does not have any effect. It just logs its result.

For more information about how to update column types in Oozie Database, see *Oozie - Updating column types in Oozie Database* section in [Step 9: Complete Post-Upgrade steps for upgrades to Cloudera Base on premises](#) .

What's New in Apache Hadoop Ozone

Learn about the new features of Apache Hadoop Ozone in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

There are no new features in this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1:

There are no new features in this release.

What's New in Apache Phoenix

Learn about the new features of Apache Phoenix in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

There are no new features in this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1:**Phoenix now supports alternate HBase connection registries**

You can now use the Phoenix thick client using additional HBase connection registries. Phoenix now supports a Zookeeper-less connection strategy using a Master Registry implementation.

Phoenix now introduces the following JDBC URL variants.

- jdbc:phoenix+zk: Uses Zookeeper. This is the original registry supported since the inception of HBase and Phoenix.
- jdbc:phoenix+rpc: Uses RPC to connect to the specified HBase Region Server or Master nodes.
- jdbc:phoenix+master: Uses RPC to connect to the specified HBase Master nodes.

For more information, see .

Phoenix FIPS support

Phoenix is now Federal Information Processing Standards (FIPS) compliant. For more information, see [Phoenix is FIPS compliant](#).

Phoenix supports rolling restart

If Phoenix service instances are running on multiple nodes, while performing a rolling restart the Phoenix services are restarted one after another ensuring zero downtime.

The OMID service also supports rolling restart and the High Availability (HA) mode for the OMID TSO server. If OMID service instances are running on multiple nodes, while performing a rolling restart the OMID services are restarted one after another ensuring zero downtime. For more information, see [Preparing for a Zero Downtime Upgrade](#).

Phoenix supports JDK 17

Phoenix supports Oracle JDK version 17.0.6 starting from CDP Runtime 7.1.9. For more information on JDK 17, see [Java Requirements](#).

Operating system support

Phoenix is now supported on the following operating systems:

- RHEL-9.1
- RHEL-8.8
- RHEL-8.8 FIPS
- Oracle-8.8 UEK
- SLES-15 SP4 for x86

What's New in Apache Ranger

Learn about the new features of Apache Ranger in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:**RAZ supports authorization for CAII (Cloudera AI Inference) service**

To support authorization for CAII, the following configurations have been added for services supported by RAZ:

- service.admin.users
- service.admin.groups

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

Cloudera Runtime 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.

What's New in Ranger KMS

Learn about the new features of Ranger KMS in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

There are no new features in this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

Cloudera Runtime 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, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1:**Enable Streams Messaging Manager principal as trusted proxy user in Schema Registry**

Streams Messaging Manager usually connects to Schema Registry on behalf of an end user. For requests coming from Streams Messaging Manager, 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, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

There are no new features in this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1:

Data Discovery and Exploration (Technical Preview)

A new Data Discovery and Exploration cluster definition is available in Cloudera 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.

- ZooKeeper SSL can now be configured for Solr and HBase Indexer. For more information, see [Enabling ZooKeeper SSL/TLS for Solr and HBase Indexer](#).
- Apache Solr is updated from 8.4.1 to 8.11.2 in this release of Cloudera Runtime. For more information, see [Apache Solr Release Notes](#) in the upstream documentation. For the list of notable unsupported features, see [Unsupported features](#).
- Using Local File System (LFS) for both MapReduce Indexer Tool (MRIT) and HBase MRIT is now supported.
- Spark 3 is now supported.
- Spark 2 is deprecated in this release and support will be dropped in an upcoming release.
- This release introduces the following two health checks for the Solr service which give information about the status of the cores hosted on different hosts:

Recovering cores

By default this check reports concerning health if any of the hosted cores are in recovering status. This threshold can be modified in the configurations with the `solr_recovering_core_thresholds` configuration parameter.

Critical cores

By default this check reports "bad health" if any of the hosted cores are in down or recovery failed status. This threshold can be modified in the configurations with the `solr_critical_core_thresholds` config.

These checks are enabled by default for the Infra Solr service but disabled by default for the Workload Solr services (Cloudera Search).

- Critical CVE fixes.

What's New in Apache Spark

Learn about the new features of Apache Spark in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

There are no new features in this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF1:

There are no new features in Spark in this release.

Cloudera Runtime 7.3.1.200 SP1:

Rebase Spark3 to Apache Spark 3.5.4 in Cloudera Runtime.

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



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

Cloudera Runtime 7.3.1.100 CHF1:

There are no new features in Spark in this release.

Cloudera Runtime 7.3.1:**Spark 3 is the default in Cloudera Runtime**

Spark 3 is the default Spark version in Cloudera Runtime. Spark 2 has been removed and no longer available in 7.3.1.0

**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](#)

What's New in Sqoop

Learn about the new features of Sqoop in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

To access the latest Sqoop documentation on Cloudera documentation website, go to [Sqoop Documentation 1.4.7.7.1.6.0](#).

Cloudera Runtime 7.3.1.500 SP3:

There are no new features in this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1:

There are no new features in this release.

What's New in Streams Messaging Manager

Learn about the new features of Streams Messaging Manager in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

There are no new features in this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

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 Streams Messaging Manager

The following security related headers were added to Streams Messaging Manager UI endpoints:

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

- Cross-Origin-Resource-Policy

Streams Messaging Manager uses trusted proxy authentication when connecting to Schema Registry

You can only interact with schemas through Streams Messaging Manager if the necessary Ranger policies are set up for Schema Registry. For Streams Messaging Manager 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 Streams Messaging Manager UI is also updated and uses the new endpoint for connector creation.

Jersey client timeout now configurable

Streams Messaging Manager 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 in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

There are no new features in this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

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 Streams Replication Manager 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 Streams Replication Manager 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
replication.records.lag.calc.enabled	true	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 Streams Replication Manager performance. If you are experiencing performance issues, you can try setting this property to false to disable the calculation of replication-records-lag. Alternatively, you can try fine-tuning how Streams Replication Manager calculates replication-records-lag with the replication.records.lag.calc.period.ms and replication.records.lag.end.offset.timeout.ms properties.
replication.records.lag.calc.period.ms	0	Controls how frequently Streams Replication Manager calculates the replication-records-lag metric. The default value of 0 means that the metric is calculated continuously. Cloudera recommends configuring this property to 15000 ms (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 Streams Replication Manager performance.

Property	Default Value	Description
replication.records.lag.end.offset.timeout.ms	6000	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 6000 ms (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 Streams Replication Manager performance.

What's New in Yarn and Yarn Queue Manager

Learn about the new features of Yarn and Yarn Queue Manager in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

There are no new features in this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

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

What's New in Apache ZooKeeper

Learn about the new features of ZooKeeper in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

There are no new features in this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1.200 SP1:

There are no new features in this release.

Cloudera Runtime 7.3.1.100 CHF 1:

There are no new features in this release.

Cloudera Runtime 7.3.1:

There are no new features in this release.

What's new in Platform Support

You must be aware of the platform support for the Cloudera Runtime 7.3.1 release, its service packs and cumulative hotfixes.

This section describes the platform support changes for the Cloudera Runtime 7.3.1 its service packs and cumulative hotfixes associated with Cloudera Base on premises 7.3.1 its service packs and cumulative hotfixes

Cloudera Runtime 7.3.1.500 SP3**New Database Support**

Postgres 17 is now supported from 7.3.1.500 SP3.

Cloudera Runtime 7.3.1.400 SP2**Platform Support Enhancements**

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

Cloudera Runtime 7.3.1.300 SP1 CHF 1**Platform Support Enhancements**

- **New OS support:**
 - RHEL 9.5

Cloudera Runtime 7.3.1.200 SP1**New OS Support**

- **Rocky Linux support for Cloudera Runtime 7.3.1.200 SP1:**
 - Rocky Linux 9.4

Starting with the Cloudera Runtime 7.3.1.200 SP1 release, Cloudera Runtime 7.3.1.200 SP1 provides support for Rocky Linux. This update ensures seamless compatibility with Rocky Linux version 9.4, offering greater flexibility and platform options.

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

Cloudera Runtime 7.3.1.100 CHF 1

There are no new Platform Support Enhancements in this release.

Cloudera Runtime 7.3.1

Platform Support Enhancements

- **New OS support:**
 - RHEL 9.4
- **New Database support:** None
- **New JDK Version:** JDK 17

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

Cloudera Runtime Component Versions

The following sections lists the the official component versions for Cloudera Runtime 7.3.1 and its Service Packs and Hotfixes. 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.

Cloudera Runtime 7.3.1.500 SP3

List of the official component versions for Cloudera Runtime 7.3.1.500 SP3.

Apache Components

Component	Version
Apache Arrow	0.11.1.7.3.1.500-182
Apache Atlas	2.1.0.7.3.1.500-182
Apache Calcite	1.25.0.7.3.1.500-182
Apache Avatica	1.22.0.7.3.1.500-182
Apache Avro	1.11.1.7.3.1.500-182
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.3.1.500-182
Apache HBase	2.4.17.7.3.1.500-182
Apache Hive	3.1.3000.7.3.1.500-182
Apache Iceberg	1.5.2.7.3.1.500-182
Apache Impala	4.4.0.7.3.1.500-182
Apache Kafka	3.4.1.7.3.1.500-182
Apache Knox	2.0.0.7.3.1.500-182
Apache Kudu	1.17.0.7.3.1.500-182
Apache Livy	0.7.23000.7.3.1.500-182
Apache MapReduce	3.1.1.7.3.1.500-182
Apache Ozone	• 7.3.1.0 - 1.3.0.7.3.1.500-182 • 7.3.1.100 and higher - 1.4.0.7.3.1.100-39

Component	Version
Apache Oozie	5.1.0.7.3.1.500-182
Apache ORC	1.8.3.7.3.1.500-182
Apache Parquet	1.12.3.7.3.1.500-182
Apache Phoenix	5.1.3.7.3.1.500-182
Apache Ranger	2.4.0.7.3.1.500-182
Apache Solr	8.11.2.7.3.1.500-182
Apache Spark	3.5.4.7.3.1.500-182
Apache Sqoop	1.4.7.7.3.1.500-182
Apache Tez	0.9.1.7.3.1.500-182
Apache ZooKeeper	3.8.1.7.3.1.500-182

Other Components

Component	Version
Cruise Control	2.5.116.7.3.1.500-182
Data Analytics Studio	1.4.2.7.3.1.500-182
GCS Connector	2.1.2.7.3.1.500-182
Hue	4.11.0.7.3.1.500-182
Search	1.0.0.7.3.1.500-182
Schema Registry	0.10.0.7.3.1.500-182
Streams Messaging Manager	2.3.0.7.3.1.500-182
Streams Replication Manager	1.1.0.7.3.1.500-182
Data Discovery and Exploration	Technical Preview

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.3.1.500-182
Hive Meta Store (HMS)	1.0.0.7.3.1.500-182
Hive on Tez	1.0.0.7.3.1.500-182
Hive Warehouse Connector	1.0.0.7.3.1.500-182
Spark Atlas Connector	3.4.1.7.3.1.500-182
Spark Schema Registry	3.4.1.7.3.1.500-182



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.
- Starting from 7.3.1.100 release, the Apache Ozone version is Apache Ozone 1.4.0.

Cloudera Runtime 7.3.1.400 SP2

List of the official component versions for Cloudera Runtime 7.3.1.400 SP2.

Apache Components

Component	Version
Apache Arrow	0.11.1.7.3.1.400-100
Apache Atlas	2.1.0.7.3.1.400-100
Apache Calcite	1.25.0.7.3.1.400-100
Apache Avatica	1.22.0.7.3.1.400-100
Apache Avro	1.11.1.7.3.1.400-100
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.3.1.400-100
Apache HBase	2.4.17.7.3.1.400-100
Apache Hive	3.1.3000.7.3.1.400-100
Apache Iceberg	1.3.1.7.3.1.400-100
Apache Impala	4.0.0.7.3.1.400-100
Apache Kafka	3.4.1.7.3.1.400-100
Apache Knox	2.0.0.7.3.1.400-100
Apache Kudu	1.17.0.7.3.1.400-100
Apache Livy	0.7.23000.7.3.1.400-100
Apache MapReduce	3.1.1.7.3.1.400-100
Apache Ozone	7.3.1.0 - 1.3.0.7.3.1.400-100 7.3.1.100 and higher - 1.4.0.7.3.1.100-39
Apache Oozie	5.1.0.7.3.1.400-100
Apache ORC	1.8.3.7.3.1.400-100
Apache Parquet	1.12.3.7.3.1.400-100
Apache Phoenix	5.1.3.7.3.1.400-100
Apache Ranger	2.4.0.7.3.1.400-100
Apache Solr	8.11.2.7.3.1.400-100
Apache Spark	3.5.4.7.3.1.400-100
Apache Sqoop	1.4.7.7.3.1.400-100
Apache Tez	0.9.1.7.3.1.400-100
Apache ZooKeeper	3.8.1.7.3.1.400-100

Other Components

Component	Version
Cruise Control	2.5.116.7.3.1.400-100
Data Analytics Studio	1.4.2.7.3.1.400-100
GCS Connector	2.1.2.7.3.1.400-100
Hue	4.11.0.7.3.1.400-100
Search	1.0.0.7.3.1.400-100
Schema Registry	0.10.0.7.3.1.400-100
Streams Messaging Manager	2.3.0.7.3.1.400-100
Streams Replication Manager	1.1.0.7.3.1.400-100
Data Discovery and Exploration	Technical Preview

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.3.1.400-100
Hive Meta Store (HMS)	1.0.0.7.3.1.400-100
Hive on Tez	1.0.0.7.3.1.400-100
Hive Warehouse Connector	1.0.0.7.3.1.400-100
Spark Atlas Connector	3.4.1.7.3.1.400-100
Spark Schema Registry	3.4.1.7.3.1.400-100

**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.
- Starting from 7.3.1.100 release, the Apache Ozone version is Apache Ozone 1.4.0.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

List of the official component versions for Cloudera Runtime 7.3.1.300 SP1 CHF 1.

Apache Components

Component	Version
Apache Arrow	0.11.1.7.3.1.300-81
Apache Atlas	2.1.0.7.3.1.300-81
Apache Calcite	1.25.0.7.3.1.300-81
Apache Avatica	1.22.0.7.3.1.300-81
Apache Avro	1.11.1.7.3.1.300-81
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.3.1.300-81
Apache HBase	2.4.17.7.3.1.300-81
Apache Hive	3.1.3000.7.3.1.300-81
Apache Iceberg	1.3.1.7.3.1.300-81
Apache Impala	4.0.0.7.3.1.300-81
Apache Kafka	3.4.1.7.3.1.300-81
Apache Knox	2.0.0.7.3.1.300-81
Apache Kudu	1.17.0.7.3.1.300-81
Apache Livy	0.7.23000.7.3.1.300-81
Apache MapReduce	3.1.1.7.3.1.300-81
Apache Ozone	• 7.3.1.0 - 1.3.0.7.3.1.300-81 • 7.3.1.100 and higher - 1.4.0.7.3.1.100-39
Apache Oozie	5.1.0.7.3.1.300-81
Apache ORC	1.8.3.7.3.1.300-81
Apache Parquet	1.12.3.7.3.1.300-81
Apache Phoenix	5.1.3.7.3.1.300-81
Apache Ranger	2.4.0.7.3.1.300-81
Apache Solr	8.11.2.7.3.1.300-81

Component	Version
Apache Spark	3.5.4.7.3.1.300-81
Apache Sqoop	1.4.7.7.3.1.300-81
Apache Tez	0.9.1.7.3.1.300-81
Apache ZooKeeper	3.8.1.7.3.1.300-81

Other Components

Component	Version
Cruise Control	2.5.116.7.3.1.300-81
Data Analytics Studio	1.4.2.7.3.1.300-81
GCS Connector	2.1.2.7.3.1.300-81
Hue	4.11.0.7.3.1.300-81
Search	1.0.0.7.3.1.300-81
Schema Registry	0.10.0.7.3.1.300-81
Streams Messaging Manager	2.3.0.7.3.1.300-81
Streams Replication Manager	1.1.0.7.3.1.300-81
Data Discovery and Exploration	Technical Preview

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.3.1.300-81
Hive Meta Store (HMS)	1.0.0.7.3.1.300-81
Hive on Tez	1.0.0.7.3.1.300-81
Hive Warehouse Connector	1.0.0.7.3.1.300-81
Spark Atlas Connector	3.4.1.7.3.1.300-81
Spark Schema Registry	3.4.1.7.3.1.300-81



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.
- Starting from 7.3.1.100 release, the Apache Ozone version is Apache Ozone 1.4.0.

Cloudera Runtime 7.3.1.200 SP1

List of the official component versions for Cloudera Runtime 7.3.1.200 SP1.

Apache Components

Component	Version
Apache Arrow	0.11.1.7.3.1.200-90
Apache Atlas	2.1.0.7.3.1.200-90
Apache Calcite	1.25.0.7.3.1.200-90
Apache Avatica	1.22.0.7.3.1.200-90
Apache Avro	1.11.1.7.3.1.200-90

Component	Version
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.3.1.200-90
Apache HBase	2.4.17.7.3.1.200-90
Apache Hive	3.1.3000.7.3.1.200-90
Apache Iceberg	1.3.1.7.3.1.200-90
Apache Impala	4.0.0.7.3.1.200-90
Apache Kafka	3.4.1.7.3.1.200-90
Apache Knox	2.0.0.7.3.1.200-90
Apache Kudu	1.17.0.7.3.1.200-90
Apache Livy	0.7.23000.7.3.1.200-90
Apache MapReduce	3.1.1.7.3.1.200-90
Apache Ozone	7.3.1.0 - 1.3.0.7.3.1.200-90 7.3.1.100 and higher - 1.4.0.7.3.1.100-39
Apache Oozie	5.1.0.7.3.1.200-90
Apache ORC	1.8.3.7.3.1.200-90
Apache Parquet	1.12.3.7.3.1.200-90
Apache Phoenix	5.1.3.7.3.1.200-90
Apache Ranger	2.4.0.7.3.1.200-90
Apache Solr	8.11.2.7.3.1.200-90
Apache Spark	3.5.4.7.3.1.200-90
Apache Sqoop	1.4.7.7.3.1.200-90
Apache Tez	0.9.1.7.3.1.200-90
Apache ZooKeeper	3.8.1.7.3.1.200-90

Other Components

Component	Version
Cruise Control	2.5.116.7.3.1.200-90
Data Analytics Studio	1.4.2.7.3.1.200-90
GCS Connector	2.1.2.7.3.1.200-90
Hue	4.11.0.7.3.1.200-90
Search	1.0.0.7.3.1.200-90
Schema Registry	0.10.0.7.3.1.200-90
Streams Messaging Manager	2.3.0.7.3.1.200-90
Streams Replication Manager	1.1.0.7.3.1.200-90
Data Discovery and Exploration	Technical Preview

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.3.1.200-90
Hive Meta Store (HMS)	1.0.0.7.3.1.200-90
Hive on Tez	1.0.0.7.3.1.200-90

Component	Version
Hive Warehouse Connector	1.0.0.7.3.1.200-90
Spark Atlas Connector	3.4.1.7.3.1.200-90
Spark Schema Registry	3.4.1.7.3.1.200-90

**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.
- Starting from 7.3.1.100 release, the Apache Ozone version is Apache Ozone 1.4.0.

Cloudera Runtime 7.3.1.100 CHF 1

List of the official component versions for Cloudera Runtime 7.3.1.100 CHF 1.

Apache Components

Component	Version
Apache Arrow	0.11.1.7.3.1.100-57
Apache Atlas	2.1.0.7.3.1.100-57
Apache Calcite	1.25.0.7.3.1.100-57
Apache Avatica	1.22.0.7.3.1.100-57
Apache Avro	1.11.1.7.3.1.100-57
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.3.1.100-57
Apache HBase	2.4.17.7.3.1.100-57
Apache Hive	3.1.3000.7.3.1.100-57
Apache Iceberg	1.3.1.7.3.1.100-57
Apache Impala	4.0.0.7.3.1.100-57
Apache Kafka	3.4.1.7.3.1.100-57
Apache Knox	2.0.0.7.3.1.100-57
Apache Kudu	1.17.0.7.3.1.100-57
Apache Livy	0.7.23000.7.3.1.100-57
Apache MapReduce	3.1.1.7.3.1.100-57
Apache Ozone	• 7.3.1.0 - 1.3.0.7.3.1.100-57 • 7.3.1.100 and higher - 1.4.0.7.3.1.100-39
Apache Oozie	5.1.0.7.3.1.100-57
Apache ORC	1.8.3.7.3.1.100-57
Apache Parquet	1.12.3.7.3.1.100-57
Apache Phoenix	5.1.3.7.3.1.100-57
Apache Ranger	2.4.0.7.3.1.100-57
Apache Solr	8.11.2.7.3.1.100-57
Apache Spark	3.4.1.7.3.1.100-57
Apache Sqoop	1.4.7.7.3.1.100-57
Apache Tez	0.9.1.7.3.1.100-57
Apache ZooKeeper	3.8.1.7.3.1.100-57

Other Components

Component	Version
Cruise Control	2.5.116.7.3.1.100-57
Data Analytics Studio	1.4.2.7.3.1.100-57
GCS Connector	2.1.2.7.3.1.100-57
Hue	4.11.0.7.3.1.100-57
Search	1.0.0.7.3.1.100-57
Schema Registry	0.10.0.7.3.1.100-57
Streams Messaging Manager	2.3.0.7.3.1.100-57
Streams Replication Manager	1.1.0.7.3.1.100-57
Data Discovery and Exploration	Technical Preview

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.3.1.100-57
Hive Meta Store (HMS)	1.0.0.7.3.1.100-57
Hive on Tez	1.0.0.7.3.1.100-57
Hive Warehouse Connector	1.0.0.7.3.1.100-57
Spark Atlas Connector	3.4.1.7.3.1.100-57
Spark Schema Registry	3.4.1.7.3.1.100-57

**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.
- Starting from 7.3.1.100 release, the Apache Ozone version is Apache Ozone 1.4.0.

Cloudera Runtime 7.3.1

List of the official component versions for Cloudera Runtime 7.3.1.

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
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

Component	Version
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	7.3.1.0 - 1.3.0.7.3.1.0-197 7.3.1.100 and higher - 1.4.0.7.3.1.100-39
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	4.11.0.7.3.1.0-197
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.
- Starting from 7.3.1.100 release, the Apache Ozone version is Apache Ozone 1.4.0.

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>
```

Cloudera Runtime 7.3.1.500-182

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	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-classification-updater	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-client-common	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-client-v1	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-client-v2	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-common	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-distro	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-docs	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-graphdb-api	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-graphdb-common	2.1.0.7.3.1.500-182

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-graphdb-janus	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-hdfs-bridge	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-index-repair-tool	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-intg	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-notification	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-plugin-classloader	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-repository	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-server-api	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	atlas-testtools	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	hbase-bridge	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	hbase-bridge-shim	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	hbase-testing-util	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	hdfs-model	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	hive-bridge	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	hive-bridge-shim	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	impala-bridge	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	impala-bridge-shim	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	impala-hook-api	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	kafka-bridge	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	kafka-bridge-shim	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	navigator-to-atlas	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	sample-app	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	sqoop-bridge	2.1.0.7.3.1.500-182
Atlas	org.apache.atlas	sqoop-bridge-shim	2.1.0.7.3.1.500-182
Avro	org.apache.avro	avro	1.11.1.7.3.1.500-182
Avro	org.apache.avro	avro-android	1.11.1.7.3.1.500-182
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.3.1.500-182
Avro	org.apache.avro	avro-compiler	1.11.1.7.3.1.500-182
Avro	org.apache.avro	avro-grpc	1.11.1.7.3.1.500-182
Avro	org.apache.avro	avro-ipc	1.11.1.7.3.1.500-182
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.3.1.500-182
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.3.1.500-182
Avro	org.apache.avro	avro-mapred	1.11.1.7.3.1.500-182
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.3.1.500-182
Avro	org.apache.avro	avro-perf	1.11.1.7.3.1.500-182
Avro	org.apache.avro	avro-protobuf	1.11.1.7.3.1.500-182
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.3.1.500-182
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.3.1.500-182

Project	groupId	artifactId	version
Avro	org.apache.avro	avro-thrift	1.11.1.7.3.1.500-182
Avro	org.apache.avro	avro-tools	1.11.1.7.3.1.500-182
Avro	org.apache.avro	trevni-avro	1.11.1.7.3.1.500-182
Avro	org.apache.avro	trevni-core	1.11.1.7.3.1.500-182
Calcite	org.apache.calcite	calcite-babel	1.25.0.7.3.1.500-182
Calcite	org.apache.calcite	calcite-core	1.25.0.7.3.1.500-182
Calcite	org.apache.calcite	calcite-druid	1.25.0.7.3.1.500-182
Calcite	org.apache.calcite	calcite-linq4j	1.25.0.7.3.1.500-182
Calcite	org.apache.calcite	calcite-plus	1.25.0.7.3.1.500-182
Calcite	org.apache.calcite	calcite-server	1.25.0.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-benchmark	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.3.1.500-182

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-miniclusterc	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.3.1.500-182

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.3.1.500-182
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.3.1.500-182
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-client	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-common	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-examples	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-http	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-it	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-logging	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-replication	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-rest	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-server	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-shell	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.3.1.500-182

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.3.1.500-182
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.3.1.500-182
Hive	org.apache.hive	catalogd-unit	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-beeline	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-classification	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-cli	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-common	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-contrib	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-exec	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-hms-catalog	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-iceberg-catalog	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-iceberg-handler	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-iceberg-shading	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-impala	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-it-iceberg	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-it-impala	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-it-util	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-jmh	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-metastore	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-parser	3.1.3000.7.3.1.500-182

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-serde	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-service	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-shims	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-streaming	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-testutils	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-udf	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	kafka-handler	3.1.3000.7.3.1.500-182
Hive	org.apache.hive	patched-iceberg-api	patched-1.5.2.7.3.1.500-182-3.1.3000
Hive	org.apache.hive	patched-iceberg-core	patched-1.5.2.7.3.1.500-182-3.1.3000
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector-spark3_2.12	1.0.0.7.3.1.500-182
Kafka	org.apache.kafka	ci	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	connect	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	connect-api	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	connect-file	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	connect-json	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	generator	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.3.1.500-182

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.3.1.500-182
Kafka	org.apache.kafka	trogdor	3.4.1.7.3.1.500-182
Knox	org.apache.knox	gateway-adapter	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-admin-ui	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-applications	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-cloud-bindings	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-demo-ldap	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-demo-ldap-launcher	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-discovery-ambari	2.0.0.7.3.1.500-182

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-discovery-cm	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-docker	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-i18n	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-i18n-logging-log4j	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-i18n-logging-sl4j	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-openapi-ui	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-performance-test	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-ha	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-identity-assertion-common	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-identity-assertion-no-doas	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-jersey	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-rewrite	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-rewrite-common	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-security-authc-anon	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-security-authz-acls	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-security-authz-composite	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-security-authz-path-acls	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-security-clientcert	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-security-hadoopauth	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-security-jwt	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-security-pac4j	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-security-preauth	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-security-shiro	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-provider-security-webappsec	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-release	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-server	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-server-launcher	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-server-xforwarded-filter	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-admin	2.0.0.7.3.1.500-182

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-service-as	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-auth	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-definitions	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-hashicorp-vault	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-hbase	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-health	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-hive	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-idbroker	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-idbroker-plugins	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-impala	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-jkg	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-knoxsso	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-knoxssout	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-knoxtoken	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-livy	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-metadata	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-nifi	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-nifi-registry	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-remoteconfig	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-rm	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-session	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-storm	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-test	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-tgs	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-vault	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-service-webhdfs	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-shell	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-shell-launcher	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-shell-release	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-shell-samples	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-spi	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-spi-common	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-test	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-test-idbroker	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-test-release-utils	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-test-utils	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-topology-hadoop-xml	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-topology-simple	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-util-common	2.0.0.7.3.1.500-182

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-util-configinjector	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-util-launcher	2.0.0.7.3.1.500-182
Knox	org.apache.knox	gateway-util-urltemplate	2.0.0.7.3.1.500-182
Knox	org.apache.knox	hadoop-examples	2.0.0.7.3.1.500-182
Knox	org.apache.knox	knox-cli-launcher	2.0.0.7.3.1.500-182
Knox	org.apache.knox	knox-homepage-ui	2.0.0.7.3.1.500-182
Knox	org.apache.knox	knox-token-generation-ui	2.0.0.7.3.1.500-182
Knox	org.apache.knox	knox-token-management-ui	2.0.0.7.3.1.500-182
Knox	org.apache.knox	knox-webshell-ui	2.0.0.7.3.1.500-182
Knox	org.apache.knox	webhdfs-kerb-test	2.0.0.7.3.1.500-182
Knox	org.apache.knox	webhdfs-test	2.0.0.7.3.1.500-182
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.3.1.500-182
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.3.1.500-182
Kudu	org.apache.kudu	kudu-client	1.17.0.7.3.1.500-182
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.3.1.500-182
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.3.1.500-182
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.3.1.500-182
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.3.1.500-182
Livy	org.apache.livy	livy-api	0.7.23000.7.3.1.500-182
Livy	org.apache.livy	livy-client-common	0.7.23000.7.3.1.500-182
Livy	org.apache.livy	livy-client-http	0.7.23000.7.3.1.500-182
Livy	org.apache.livy	livy-core_2.12	0.7.23000.7.3.1.500-182
Livy	org.apache.livy	livy-examples	0.7.23000.7.3.1.500-182
Livy	org.apache.livy	livy-integration-test	0.7.23000.7.3.1.500-182
Livy	org.apache.livy	livy-repl_2.12	0.7.23000.7.3.1.500-182
Livy	org.apache.livy	livy-rsc	0.7.23000.7.3.1.500-182
Livy	org.apache.livy	livy-scala-api_2.12	0.7.23000.7.3.1.500-182
Livy	org.apache.livy	livy-server	0.7.23000.7.3.1.500-182
Livy	org.apache.livy	livy-test-lib	0.7.23000.7.3.1.500-182
Livy	org.apache.livy	livy-thriftserver	0.7.23000.7.3.1.500-182
Livy	org.apache.livy	livy-thriftserver-session	0.7.23000.7.3.1.500-182
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-analyzers-openssl	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.3.1.500-182

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-core	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-join	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.3.1.500-182
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.3.1.500-182
Oozie	org.apache.oozie	oozie-client	5.1.0.7.3.1.500-182
Oozie	org.apache.oozie	oozie-core	5.1.0.7.3.1.500-182
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.3.1.500-182
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.3.1.500-182
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.3.1.500-182
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.3.1.500-182
Oozie	org.apache.oozie	oozie-server	5.1.0.7.3.1.500-182
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.3.1.500-182
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.3.1.500-182
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.3.1.500-182
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.3.1.500-182
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.3.1.500-182
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.3.1.500-182
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.3.1.500-182
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.3.1.500-182
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.3.1.500-182

Project	groupId	artifactId	version
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.3.1.500-182
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.3.1.500-182
ORC	org.apache.orc	orc-core	1.8.3.7.3.1.500-182
ORC	org.apache.orc	orc-examples	1.8.3.7.3.1.500-182
ORC	org.apache.orc	orc-mapreduce	1.8.3.7.3.1.500-182
ORC	org.apache.orc	orc-shims	1.8.3.7.3.1.500-182
ORC	org.apache.orc	orc-tools	1.8.3.7.3.1.500-182
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-client	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-common	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-config	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-crypto-api	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-crypto-default	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-managed-rocksdbs	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-client	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-common	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-filesystem	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-filesystem-common	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-filesystem-hadoop2	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-filesystem-hadoop3	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-filesystem-shaded	1.4.0.7.3.1.500-182

Project	groupId	artifactId	version
Ozone	org.apache.ozone	ozone-httpfsgateway	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.3.1.500-182
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.3.1.500-182
Parquet	org.apache.parquet	parquet-avro	1.12.3.7.3.1.500-182
Parquet	org.apache.parquet	parquet-cascading	1.12.3.7.3.1.500-182
Parquet	org.apache.parquet	parquet-cascading3	1.12.3.7.3.1.500-182
Parquet	org.apache.parquet	parquet-column	1.12.3.7.3.1.500-182
Parquet	org.apache.parquet	parquet-common	1.12.3.7.3.1.500-182
Parquet	org.apache.parquet	parquet-encoding	1.12.3.7.3.1.500-182
Parquet	org.apache.parquet	parquet-format-structures	1.12.3.7.3.1.500-182
Parquet	org.apache.parquet	parquet-generator	1.12.3.7.3.1.500-182
Parquet	org.apache.parquet	parquet-hadoop	1.12.3.7.3.1.500-182
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.12.3.7.3.1.500-182
Parquet	org.apache.parquet	parquet-jackson	1.12.3.7.3.1.500-182
Parquet	org.apache.parquet	parquet-pig	1.12.3.7.3.1.500-182
Parquet	org.apache.parquet	parquet-pig-bundle	1.12.3.7.3.1.500-182
Parquet	org.apache.parquet	parquet-protobuf	1.12.3.7.3.1.500-182
Parquet	org.apache.parquet	parquet-scala_2.12	1.12.3.7.3.1.500-182
Parquet	org.apache.parquet	parquet-thrift	1.12.3.7.3.1.500-182
Parquet	org.apache.parquet	parquet-tools	1.12.3.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.3.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.3.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix-core	5.1.3.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.1.6	5.1.3.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.2.5	5.1.3.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.3.0	5.1.3.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.0	5.1.3.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.1.3.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.0	5.1.3.7.3.1.500-182

Project	groupId	artifactId	version
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.4	5.1.3.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.3.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.3.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.3.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix5-spark3	6.0.0.7.3.1.500-182
Phoenix	org.apache.phoenix	phoenix5-spark3-shaded	6.0.0.7.3.1.500-182
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	jisql	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-caii-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-gs-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.3.1.500-182

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-raz-hook-s3	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-raz-s3	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-raz-s3-lib	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-s3-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.3.1.500-182

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-trino-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-util	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	sample-client	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	sampleapp	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	shaded-raz-hook-s3	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.3.1.500-182
Ranger	org.apache.ranger	unixusersync	2.4.0.7.3.1.500-182
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.3.1.500-182
Solr	org.apache.solr	solr-analytics	8.11.2.7.3.1.500-182
Solr	org.apache.solr	solr-cell	8.11.2.7.3.1.500-182
Solr	org.apache.solr	solr-core	8.11.2.7.3.1.500-182
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.3.1.500-182
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.3.1.500-182
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.3.1.500-182
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.3.1.500-182
Solr	org.apache.solr	solr-langid	8.11.2.7.3.1.500-182
Solr	org.apache.solr	solr-ltr	8.11.2.7.3.1.500-182
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.3.1.500-182
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.3.1.500-182
Solr	org.apache.solr	solr-security-util	8.11.2.7.3.1.500-182
Solr	org.apache.solr	solr-solrj	8.11.2.7.3.1.500-182
Solr	org.apache.solr	solr-test-framework	8.11.2.7.3.1.500-182
Solr	org.apache.solr	solr-velocity	8.11.2.7.3.1.500-182
Spark	org.apache.spark	spark-avro_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-catalyst_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-common-utils_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-connect-client-jvm_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-connect-common_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-connect_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-core_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-graphx_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-hadoop-cloud_2.12	3.5.4.7.3.1.500-182

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-hive_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-kubernetes_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-kvstore_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-launcher_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-mllib-local_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-mllib_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-network-common_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-network-shuffle_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-network-yarn_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-protobuf_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-repl_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-shaded-raz	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-sketch_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-sql-api_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-sql-kafka-0-10_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-sql_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-streaming_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-tags_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-unsafe_2.12	3.5.4.7.3.1.500-182
Spark	org.apache.spark	spark-yarn_2.12	3.5.4.7.3.1.500-182
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.3.1.500-182
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.3.1.500-182
Tez	org.apache.tez	hadoop-shim	0.9.1.7.3.1.500-182
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.3.1.500-182
Tez	org.apache.tez	tez-api	0.9.1.7.3.1.500-182
Tez	org.apache.tez	tez-aux-services	0.9.1.7.3.1.500-182
Tez	org.apache.tez	tez-common	0.9.1.7.3.1.500-182
Tez	org.apache.tez	tez-dag	0.9.1.7.3.1.500-182
Tez	org.apache.tez	tez-examples	0.9.1.7.3.1.500-182
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.3.1.500-182
Tez	org.apache.tez	tez-history-parser	0.9.1.7.3.1.500-182
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.3.1.500-182
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.3.1.500-182
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.3.1.500-182
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.3.1.500-182
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.3.1.500-182

Project	groupId	artifactId	version
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.3.1.500-182
Tez	org.apache.tez	tez-tests	0.9.1.7.3.1.500-182
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.3.1.500-182
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.3.1.500-182
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.3.1.500-182
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.3.1.500-182
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.3.1.500-182
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.3.1.500-182
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.3.1.500-182
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.3.1.500-182
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.3.1.500-182
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.3.1.500-182
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.3.1.500-182
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.3.1.500-182
Zeppelin	org.apache.zeppelin	zeppelin-shaded-raz	0.8.2.7.3.1.500-182
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.3.1.500-182
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.3.1.500-182
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.3.1.500-182
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.3.1.500-182
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.3.1.500-182
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.3.1.500-182
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.3.1.500-182
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.3.1.500-182
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.3.1.500-182
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.3.1.500-182
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.3.1.500-182
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.3.1.500-182
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.3.1.500-182

Cloudera Runtime 7.3.1.400-100

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	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-classification-updater	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-client-common	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-client-v1	2.1.0.7.3.1.400-100

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-client-v2	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-common	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-distro	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-docs	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-graphdb-api	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-graphdb-common	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-graphdb-janus	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-hdfs-bridge	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-index-repair-tool	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-intg	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-notification	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-plugin-classloader	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-repository	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-server-api	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	atlas-testtools	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	hbase-bridge	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	hbase-bridge-shim	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	hbase-testing-util	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	hdfs-model	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	hive-bridge	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	hive-bridge-shim	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	impala-bridge	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	impala-bridge-shim	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	impala-hook-api	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	kafka-bridge	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	kafka-bridge-shim	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	navigator-to-atlas	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	sample-app	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	sqoop-bridge	2.1.0.7.3.1.400-100
Atlas	org.apache.atlas	sqoop-bridge-shim	2.1.0.7.3.1.400-100
Avro	org.apache.avro	avro	1.11.1.7.3.1.400-100
Avro	org.apache.avro	avro-android	1.11.1.7.3.1.400-100
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.3.1.400-100
Avro	org.apache.avro	avro-compiler	1.11.1.7.3.1.400-100
Avro	org.apache.avro	avro-grpc	1.11.1.7.3.1.400-100
Avro	org.apache.avro	avro-ipc	1.11.1.7.3.1.400-100
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.3.1.400-100

Project	groupId	artifactId	version
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.3.1.400-100
Avro	org.apache.avro	avro-mapred	1.11.1.7.3.1.400-100
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.3.1.400-100
Avro	org.apache.avro	avro-perf	1.11.1.7.3.1.400-100
Avro	org.apache.avro	avro-protobuf	1.11.1.7.3.1.400-100
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.3.1.400-100
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.3.1.400-100
Avro	org.apache.avro	avro-thrift	1.11.1.7.3.1.400-100
Avro	org.apache.avro	avro-tools	1.11.1.7.3.1.400-100
Avro	org.apache.avro	trevni-avro	1.11.1.7.3.1.400-100
Avro	org.apache.avro	trevni-core	1.11.1.7.3.1.400-100
Calcite	org.apache.calcite	calcite-babel	1.25.0.7.3.1.400-100
Calcite	org.apache.calcite	calcite-core	1.25.0.7.3.1.400-100
Calcite	org.apache.calcite	calcite-druid	1.25.0.7.3.1.400-100
Calcite	org.apache.calcite	calcite-linq4j	1.25.0.7.3.1.400-100
Calcite	org.apache.calcite	calcite-server	1.25.0.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-benchmark	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.3.1.400-100

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-minicluster	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.3.1.400-100

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.3.1.400-100
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.3.1.400-100
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-client	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-common	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-examples	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-http	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-it	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-logging	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-replication	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-rest	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-server	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.3.1.400-100

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-shell	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.3.1.400-100
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.3.1.400-100
Hive	org.apache.hive	catalogd-unit	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-beeline	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-classification	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-cli	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-common	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-contrib	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-exec	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-hms-catalog	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-iceberg-catalog	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-iceberg-handler	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-iceberg-shading	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-impala	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-it-iceberg	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-it-impala	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-it-util	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-jmh	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.3.1.400-100

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-metastore	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-parser	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-serde	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-service	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-shims	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-streaming	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-testutils	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-udf	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	kafka-handler	3.1.3000.7.3.1.400-100
Hive	org.apache.hive	patched-iceberg-api	patched-1.3.1.7.3.1.400-100-3.1.3000
Hive	org.apache.hive	patched-iceberg-core	patched-1.3.1.7.3.1.400-100-3.1.3000
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector-spark3_2.12	1.0.0.7.3.1.400-100
Kafka	org.apache.kafka	ci	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	connect	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	connect-api	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	connect-file	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	connect-json	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	generator	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.3.1.400-100

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.3.1.400-100
Kafka	org.apache.kafka	trogdor	3.4.1.7.3.1.400-100
Knox	org.apache.knox	gateway-adapter	2.0.0.7.3.1.400-100

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-admin-ui	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-applications	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-cloud-bindings	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-demo-ldap	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-demo-ldap-launcher	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-discovery-ambari	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-discovery-cm	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-docker	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-i18n	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-i18n-logging-log4j	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-i18n-logging-sl4j	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-openapi-ui	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-performance-test	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-ha	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-identity-assertion-common	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-identity-assertion-no-does	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-jersey	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-rewrite	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-rewrite-common	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-security-authc-anon	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-security-authz-acls	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-security-authz-composite	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-security-authz-path-acls	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-security-clientcert	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-security-hadoopauth	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-security-jwt	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-security-pac4j	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-security-preauth	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-provider-security-shiro	2.0.0.7.3.1.400-100

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-provider-security-webappsec	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-release	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-server	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-server-launcher	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-server-xforwarded-filter	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-admin	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-as	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-auth	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-definitions	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-hashicorp-vault	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-hbase	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-health	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-hive	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-idbroker	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-idbroker-plugins	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-impala	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-jkg	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-knoxsso	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-knoxssout	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-knoxtoken	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-livy	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-metadata	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-nifi	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-nifi-registry	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-remoteconfig	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-rm	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-session	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-storm	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-test	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-tgs	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-vault	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-service-webhdfs	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-shell	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-shell-launcher	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-shell-release	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-shell-samples	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-spi	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-spi-common	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-test	2.0.0.7.3.1.400-100

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-test-idbroker	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-test-release-utils	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-test-utils	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-topology-hadoop-xml	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-topology-simple	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-util-common	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-util-configinjector	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-util-launcher	2.0.0.7.3.1.400-100
Knox	org.apache.knox	gateway-util-urltemplate	2.0.0.7.3.1.400-100
Knox	org.apache.knox	webhdfs-test	2.0.0.7.3.1.400-100
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.3.1.400-100
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.3.1.400-100
Kudu	org.apache.kudu	kudu-client	1.17.0.7.3.1.400-100
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.3.1.400-100
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.3.1.400-100
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.3.1.400-100
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.3.1.400-100
Livy	org.apache.livy	livy-api	0.7.23000.7.3.1.400-100
Livy	org.apache.livy	livy-client-common	0.7.23000.7.3.1.400-100
Livy	org.apache.livy	livy-client-http	0.7.23000.7.3.1.400-100
Livy	org.apache.livy	livy-core_2.12	0.7.23000.7.3.1.400-100
Livy	org.apache.livy	livy-examples	0.7.23000.7.3.1.400-100
Livy	org.apache.livy	livy-integration-test	0.7.23000.7.3.1.400-100
Livy	org.apache.livy	livy-repl_2.12	0.7.23000.7.3.1.400-100
Livy	org.apache.livy	livy-rsc	0.7.23000.7.3.1.400-100
Livy	org.apache.livy	livy-scala-api_2.12	0.7.23000.7.3.1.400-100
Livy	org.apache.livy	livy-server	0.7.23000.7.3.1.400-100
Livy	org.apache.livy	livy-test-lib	0.7.23000.7.3.1.400-100
Livy	org.apache.livy	livy-thriftserver	0.7.23000.7.3.1.400-100
Livy	org.apache.livy	livy-thriftserver-session	0.7.23000.7.3.1.400-100
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-analyzers-openssl	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.3.1.400-100

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-core	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-join	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.3.1.400-100
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.3.1.400-100
Oozie	org.apache.oozie	oozie-client	5.1.0.7.3.1.400-100
Oozie	org.apache.oozie	oozie-core	5.1.0.7.3.1.400-100
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.3.1.400-100
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.3.1.400-100
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.3.1.400-100
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.3.1.400-100
Oozie	org.apache.oozie	oozie-server	5.1.0.7.3.1.400-100
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.3.1.400-100
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.3.1.400-100
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.3.1.400-100
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.3.1.400-100
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.3.1.400-100
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.3.1.400-100
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.3.1.400-100
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.3.1.400-100
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.3.1.400-100
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.3.1.400-100

Project	groupId	artifactId	version
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.3.1.400-100
ORC	org.apache.orc	orc-core	1.8.3.7.3.1.400-100
ORC	org.apache.orc	orc-examples	1.8.3.7.3.1.400-100
ORC	org.apache.orc	orc-mapreduce	1.8.3.7.3.1.400-100
ORC	org.apache.orc	orc-shims	1.8.3.7.3.1.400-100
ORC	org.apache.orc	orc-tools	1.8.3.7.3.1.400-100
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-client	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-common	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-config	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-crypto-api	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-crypto-default	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-managed-rocksdb	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-client	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-common	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-filessystem	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-filessystem-common	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-filessystem-hadoop2	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-filessystem-hadoop3	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-filessystem-shaded	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-httpsgateway	1.4.0.7.3.1.400-100

Project	groupId	artifactId	version
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.3.1.400-100
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.3.1.400-100
Parquet	org.apache.parquet	parquet-avro	1.12.3.7.3.1.400-100
Parquet	org.apache.parquet	parquet-cascading	1.12.3.7.3.1.400-100
Parquet	org.apache.parquet	parquet-cascading3	1.12.3.7.3.1.400-100
Parquet	org.apache.parquet	parquet-column	1.12.3.7.3.1.400-100
Parquet	org.apache.parquet	parquet-common	1.12.3.7.3.1.400-100
Parquet	org.apache.parquet	parquet-encoding	1.12.3.7.3.1.400-100
Parquet	org.apache.parquet	parquet-format-structures	1.12.3.7.3.1.400-100
Parquet	org.apache.parquet	parquet-generator	1.12.3.7.3.1.400-100
Parquet	org.apache.parquet	parquet-hadoop	1.12.3.7.3.1.400-100
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.12.3.7.3.1.400-100
Parquet	org.apache.parquet	parquet-jackson	1.12.3.7.3.1.400-100
Parquet	org.apache.parquet	parquet-pig	1.12.3.7.3.1.400-100
Parquet	org.apache.parquet	parquet-pig-bundle	1.12.3.7.3.1.400-100
Parquet	org.apache.parquet	parquet-protobuf	1.12.3.7.3.1.400-100
Parquet	org.apache.parquet	parquet-scala_2.12	1.12.3.7.3.1.400-100
Parquet	org.apache.parquet	parquet-thrift	1.12.3.7.3.1.400-100
Parquet	org.apache.parquet	parquet-tools	1.12.3.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.3.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.3.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix-core	5.1.3.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.1.6	5.1.3.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.2.5	5.1.3.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.3.0	5.1.3.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.0	5.1.3.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.1.3.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.0	5.1.3.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.4	5.1.3.7.3.1.400-100

Project	groupId	artifactId	version
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.3.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.3.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.3.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix5-spark3	6.0.0.7.3.1.400-100
Phoenix	org.apache.phoenix	phoenix5-spark3-shaded	6.0.0.7.3.1.400-100
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	jisql	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-gs-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.3.1.400-100

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-raz-hook-s3	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-raz-s3	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-raz-s3-lib	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-s3-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.3.1.400-100

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-trino-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-util	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	sample-client	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	sampleapp	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	shaded-raz-hook-s3	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.3.1.400-100
Ranger	org.apache.ranger	unixusersync	2.4.0.7.3.1.400-100
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.3.1.400-100
Solr	org.apache.solr	solr-analytics	8.11.2.7.3.1.400-100
Solr	org.apache.solr	solr-cell	8.11.2.7.3.1.400-100
Solr	org.apache.solr	solr-core	8.11.2.7.3.1.400-100
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.3.1.400-100
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.3.1.400-100
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.3.1.400-100
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.3.1.400-100
Solr	org.apache.solr	solr-langid	8.11.2.7.3.1.400-100
Solr	org.apache.solr	solr-ltr	8.11.2.7.3.1.400-100
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.3.1.400-100
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.3.1.400-100
Solr	org.apache.solr	solr-security-util	8.11.2.7.3.1.400-100
Solr	org.apache.solr	solr-solrj	8.11.2.7.3.1.400-100
Solr	org.apache.solr	solr-test-framework	8.11.2.7.3.1.400-100
Solr	org.apache.solr	solr-velocity	8.11.2.7.3.1.400-100
Spark	org.apache.spark	spark-avro_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-catalyst_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-common-utils_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-connect-client-jvm_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-connect-common_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-connect_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-core_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-graphx_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-hadoop-cloud_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-hive_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-kubernetes_2.12	3.5.4.7.3.1.400-100

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-kvstore_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-launcher_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-mllib-local_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-mllib_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-network-common_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-network-shuffle_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-network-yarn_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-protobuf_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-repl_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-shaded-raz	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-sketch_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-sql-api_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-sql-kafka-0-10_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-sql_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-streaming_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-tags_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-unsafe_2.12	3.5.4.7.3.1.400-100
Spark	org.apache.spark	spark-yarn_2.12	3.5.4.7.3.1.400-100
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.3.1.400-100
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.3.1.400-100
Tez	org.apache.tez	hadoop-shim	0.9.1.7.3.1.400-100
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.3.1.400-100
Tez	org.apache.tez	tez-api	0.9.1.7.3.1.400-100
Tez	org.apache.tez	tez-aux-services	0.9.1.7.3.1.400-100
Tez	org.apache.tez	tez-common	0.9.1.7.3.1.400-100
Tez	org.apache.tez	tez-dag	0.9.1.7.3.1.400-100
Tez	org.apache.tez	tez-examples	0.9.1.7.3.1.400-100
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.3.1.400-100
Tez	org.apache.tez	tez-history-parser	0.9.1.7.3.1.400-100
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.3.1.400-100
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.3.1.400-100
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.3.1.400-100
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.3.1.400-100
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.3.1.400-100
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.3.1.400-100
Tez	org.apache.tez	tez-tests	0.9.1.7.3.1.400-100

Project	groupId	artifactId	version
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.3.1.400-100
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.3.1.400-100
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.3.1.400-100
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.3.1.400-100
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.3.1.400-100
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.3.1.400-100
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.3.1.400-100
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.3.1.400-100
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.3.1.400-100
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.3.1.400-100
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.3.1.400-100
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.3.1.400-100
Zeppelin	org.apache.zeppelin	zeppelin-shaded-raz	0.8.2.7.3.1.400-100
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.3.1.400-100
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.3.1.400-100
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.3.1.400-100
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.3.1.400-100
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.3.1.400-100
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.3.1.400-100
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.3.1.400-100
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.3.1.400-100
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.3.1.400-100
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.3.1.400-100
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.3.1.400-100
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.3.1.400-100
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.3.1.400-100

Cloudera Runtime 7.3.1.300-81

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	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-classification-updater	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-client-common	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-client-v1	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-client-v2	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.1.0.7.3.1.300-81

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-common	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-distro	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-docs	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-graphdb-api	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-graphdb-common	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-graphdb-janus	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-hdfs-bridge	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-index-repair-tool	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-intg	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-notification	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-plugin-classloader	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-repository	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-server-api	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	atlas-testtools	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	hbase-bridge	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	hbase-bridge-shim	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	hbase-testing-util	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	hdfs-model	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	hive-bridge	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	hive-bridge-shim	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	impala-bridge	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	impala-bridge-shim	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	impala-hook-api	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	kafka-bridge	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	kafka-bridge-shim	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	navigator-to-atlas	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	sample-app	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	sqoop-bridge	2.1.0.7.3.1.300-81
Atlas	org.apache.atlas	sqoop-bridge-shim	2.1.0.7.3.1.300-81
Avro	org.apache.avro	avro	1.11.1.7.3.1.300-81
Avro	org.apache.avro	avro-android	1.11.1.7.3.1.300-81
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.3.1.300-81
Avro	org.apache.avro	avro-compiler	1.11.1.7.3.1.300-81
Avro	org.apache.avro	avro-grpc	1.11.1.7.3.1.300-81
Avro	org.apache.avro	avro-ipc	1.11.1.7.3.1.300-81
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.3.1.300-81
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.3.1.300-81
Avro	org.apache.avro	avro-mapred	1.11.1.7.3.1.300-81

Project	groupId	artifactId	version
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.3.1.300-81
Avro	org.apache.avro	avro-perf	1.11.1.7.3.1.300-81
Avro	org.apache.avro	avro-protobuf	1.11.1.7.3.1.300-81
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.3.1.300-81
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.3.1.300-81
Avro	org.apache.avro	avro-thrift	1.11.1.7.3.1.300-81
Avro	org.apache.avro	avro-tools	1.11.1.7.3.1.300-81
Avro	org.apache.avro	trevni-avro	1.11.1.7.3.1.300-81
Avro	org.apache.avro	trevni-core	1.11.1.7.3.1.300-81
Calcite	org.apache.calcite	calcite-babel	1.25.0.7.3.1.300-81
Calcite	org.apache.calcite	calcite-core	1.25.0.7.3.1.300-81
Calcite	org.apache.calcite	calcite-druid	1.25.0.7.3.1.300-81
Calcite	org.apache.calcite	calcite-ling4j	1.25.0.7.3.1.300-81
Calcite	org.apache.calcite	calcite-server	1.25.0.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-benchmark	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.3.1.300-81

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-miniclusterc	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.3.1.300-81

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.3.1.300-81
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.3.1.300-81
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-client	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-common	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-examples	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-http	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-it	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-logging	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-replication	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-rest	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-server	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.3.1.300-81

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-shell	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.3.1.300-81
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.3.1.300-81
Hive	org.apache.hive	catalogd-unit	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-beeline	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-classification	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-cli	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-common	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-contrib	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-exec	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-iceberg-catalog	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-iceberg-handler	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-iceberg-shading	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-impala	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-it-iceberg	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-it-impala	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-it-util	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-jmh	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.3.1.300-81

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-metastore	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-parser	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-serde	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-service	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-shims	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-streaming	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-testutils	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-udf	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	kafka-handler	3.1.3000.7.3.1.300-81
Hive	org.apache.hive	patched-iceberg-api	patched-1.3.1.7.3.1.300-81-3.1.3000.
Hive	org.apache.hive	patched-iceberg-core	patched-1.3.1.7.3.1.300-81-3.1.3000.
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector-spark3_2.12	1.0.0.7.3.1.300-81
Kafka	org.apache.kafka	ci	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	connect	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	connect-api	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	connect-file	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	connect-json	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	generator	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.3.1.300-81

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.3.1.300-81
Kafka	org.apache.kafka	trogdor	3.4.1.7.3.1.300-81
Knox	org.apache.knox	gateway-adapter	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-admin-ui	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-applications	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-cloud-bindings	2.0.0.7.3.1.300-81

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-demo-ldap	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-demo-ldap-launcher	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-discovery-ambari	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-discovery-cm	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-docker	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-i18n	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-i18n-logging-log4j	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-i18n-logging-slf4j	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-openapi-ui	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-performance-test	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-ha	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-identity-assertion-common	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-identity-assertion-no-does	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-jersey	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-rewrite	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-rewrite-common	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-security-authc-anon	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-security-authz-acls	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-security-authz-composite	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-security-authz-path-acls	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-security-clientcert	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-security-hadoopauth	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-security-jwt	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-security-pac4j	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-security-preauth	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-security-shiro	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-provider-security-webappsec	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-release	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-server	2.0.0.7.3.1.300-81

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-server-launcher	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-server-xforwarded-filter	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-admin	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-as	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-auth	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-definitions	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-hashicorp-vault	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-hbase	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-health	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-hive	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-idbroker	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-idbroker-plugins	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-impala	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-jkg	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-knoxsso	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-knoxsout	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-knoxtoken	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-livy	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-metadata	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-nifi	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-nifi-registry	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-remoteconfig	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-rm	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-session	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-storm	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-test	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-tgs	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-vault	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-service-webhdfs	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-shell	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-shell-launcher	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-shell-release	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-shell-samples	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-spi	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-spi-common	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-test	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-test-idbroker	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-test-release-utils	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-test-utils	2.0.0.7.3.1.300-81

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-topology-hadoop-xml	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-topology-simple	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-util-common	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-util-configinjector	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-util-launcher	2.0.0.7.3.1.300-81
Knox	org.apache.knox	gateway-util-urltemplate	2.0.0.7.3.1.300-81
Knox	org.apache.knox	hadoop-examples	2.0.0.7.3.1.300-81
Knox	org.apache.knox	knox-cli-launcher	2.0.0.7.3.1.300-81
Knox	org.apache.knox	knox-homepage-ui	2.0.0.7.3.1.300-81
Knox	org.apache.knox	knox-token-generation-ui	2.0.0.7.3.1.300-81
Knox	org.apache.knox	knox-token-management-ui	2.0.0.7.3.1.300-81
Knox	org.apache.knox	knox-webshell-ui	2.0.0.7.3.1.300-81
Knox	org.apache.knox	webhdfs-kerb-test	2.0.0.7.3.1.300-81
Knox	org.apache.knox	webhdfs-test	2.0.0.7.3.1.300-81
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.3.1.300-81
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.3.1.300-81
Kudu	org.apache.kudu	kudu-client	1.17.0.7.3.1.300-81
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.3.1.300-81
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.3.1.300-81
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.3.1.300-81
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.3.1.300-81
Livy	org.apache.livy	livy-api	0.7.23000.7.3.1.300-81
Livy	org.apache.livy	livy-client-common	0.7.23000.7.3.1.300-81
Livy	org.apache.livy	livy-client-http	0.7.23000.7.3.1.300-81
Livy	org.apache.livy	livy-core_2.12	0.7.23000.7.3.1.300-81
Livy	org.apache.livy	livy-examples	0.7.23000.7.3.1.300-81
Livy	org.apache.livy	livy-integration-test	0.7.23000.7.3.1.300-81
Livy	org.apache.livy	livy-repl_2.12	0.7.23000.7.3.1.300-81
Livy	org.apache.livy	livy-rsc	0.7.23000.7.3.1.300-81
Livy	org.apache.livy	livy-scala-api_2.12	0.7.23000.7.3.1.300-81
Livy	org.apache.livy	livy-server	0.7.23000.7.3.1.300-81
Livy	org.apache.livy	livy-test-lib	0.7.23000.7.3.1.300-81
Livy	org.apache.livy	livy-thriftserver	0.7.23000.7.3.1.300-81
Livy	org.apache.livy	livy-thriftserver-session	0.7.23000.7.3.1.300-81
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.3.1.300-81

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-analyzers-openssl	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-core	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-join	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.3.1.300-81
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.3.1.300-81
Oozie	org.apache.oozie	oozie-client	5.1.0.7.3.1.300-81
Oozie	org.apache.oozie	oozie-core	5.1.0.7.3.1.300-81
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.3.1.300-81
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.3.1.300-81
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.3.1.300-81
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.3.1.300-81
Oozie	org.apache.oozie	oozie-server	5.1.0.7.3.1.300-81
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.3.1.300-81
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.3.1.300-81
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.3.1.300-81
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.3.1.300-81
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.3.1.300-81
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.3.1.300-81

Project	groupId	artifactId	version
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.3.1.300-81
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.3.1.300-81
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.3.1.300-81
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.3.1.300-81
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.3.1.300-81
ORC	org.apache.orc	orc-core	1.8.3.7.3.1.300-81
ORC	org.apache.orc	orc-examples	1.8.3.7.3.1.300-81
ORC	org.apache.orc	orc-mapreduce	1.8.3.7.3.1.300-81
ORC	org.apache.orc	orc-shims	1.8.3.7.3.1.300-81
ORC	org.apache.orc	orc-tools	1.8.3.7.3.1.300-81
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-client	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-common	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-config	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-crypto-api	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-crypto-default	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-managed-rocksdbs	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-client	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-common	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-filessystem	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-filessystem-common	1.4.0.7.3.1.300-81

Project	groupId	artifactId	version
Ozone	org.apache.ozone	ozone-file-system-hadoop2	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-file-system-hadoop3	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-file-system-shaded	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-httpfs-gateway	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-recon-codegen	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-s3-gateway	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.3.1.300-81
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.3.1.300-81
Parquet	org.apache.parquet	parquet-avro	1.12.3.7.3.1.300-81
Parquet	org.apache.parquet	parquet-cascading	1.12.3.7.3.1.300-81
Parquet	org.apache.parquet	parquet-cascading3	1.12.3.7.3.1.300-81
Parquet	org.apache.parquet	parquet-column	1.12.3.7.3.1.300-81
Parquet	org.apache.parquet	parquet-common	1.12.3.7.3.1.300-81
Parquet	org.apache.parquet	parquet-encoding	1.12.3.7.3.1.300-81
Parquet	org.apache.parquet	parquet-format-structures	1.12.3.7.3.1.300-81
Parquet	org.apache.parquet	parquet-generator	1.12.3.7.3.1.300-81
Parquet	org.apache.parquet	parquet-hadoop	1.12.3.7.3.1.300-81
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.12.3.7.3.1.300-81
Parquet	org.apache.parquet	parquet-jackson	1.12.3.7.3.1.300-81
Parquet	org.apache.parquet	parquet-pig	1.12.3.7.3.1.300-81
Parquet	org.apache.parquet	parquet-pig-bundle	1.12.3.7.3.1.300-81
Parquet	org.apache.parquet	parquet-protobuf	1.12.3.7.3.1.300-81
Parquet	org.apache.parquet	parquet-scala_2.12	1.12.3.7.3.1.300-81
Parquet	org.apache.parquet	parquet-thrift	1.12.3.7.3.1.300-81
Parquet	org.apache.parquet	parquet-tools	1.12.3.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.3.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.3.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix-core	5.1.3.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.1.6	5.1.3.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.2.5	5.1.3.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.3.0	5.1.3.7.3.1.300-81

Project	groupId	artifactId	version
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.0	5.1.3.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.1.3.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.0	5.1.3.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.4	5.1.3.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.3.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.3.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.3.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix5-spark3	6.0.0.7.3.1.300-81
Phoenix	org.apache.phoenix	phoenix5-spark3-shaded	6.0.0.7.3.1.300-81
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	jisql	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-gs-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.3.1.300-81

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-raz-hook-s3	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-raz-s3	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-raz-s3-lib	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-s3-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.3.1.300-81

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-trino-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-util	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	sample-client	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	sampleapp	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	shaded-raz-hook-s3	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.3.1.300-81
Ranger	org.apache.ranger	unixusersync	2.4.0.7.3.1.300-81
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.3.1.300-81
Solr	org.apache.solr	solr-analytics	8.11.2.7.3.1.300-81
Solr	org.apache.solr	solr-cell	8.11.2.7.3.1.300-81
Solr	org.apache.solr	solr-core	8.11.2.7.3.1.300-81
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.3.1.300-81
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.3.1.300-81
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.3.1.300-81
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.3.1.300-81
Solr	org.apache.solr	solr-langid	8.11.2.7.3.1.300-81
Solr	org.apache.solr	solr-ltr	8.11.2.7.3.1.300-81
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.3.1.300-81
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.3.1.300-81
Solr	org.apache.solr	solr-security-util	8.11.2.7.3.1.300-81
Solr	org.apache.solr	solr-solrj	8.11.2.7.3.1.300-81
Solr	org.apache.solr	solr-test-framework	8.11.2.7.3.1.300-81
Solr	org.apache.solr	solr-velocity	8.11.2.7.3.1.300-81
Spark	org.apache.spark	spark-avro_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-catalyst_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-common-utils_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-connect-client-jvm_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-connect-common_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-connect_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-core_2.12	3.5.4.7.3.1.300-81

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-graphx_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-hadoop-cloud_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-hive_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-kubernetes_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-kvstore_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-launcher_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-mllib-local_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-mllib_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-network-common_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-network-shuffle_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-network-yarn_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-protobuf_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-repl_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-shaded-raz	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-sketch_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-sql-api_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-sql-kafka-0-10_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-sql_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-streaming_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-tags_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-unsafe_2.12	3.5.4.7.3.1.300-81
Spark	org.apache.spark	spark-yarn_2.12	3.5.4.7.3.1.300-81
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.3.1.300-81
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.3.1.300-81
Tez	org.apache.tez	hadoop-shim	0.9.1.7.3.1.300-81
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.3.1.300-81
Tez	org.apache.tez	tez-api	0.9.1.7.3.1.300-81
Tez	org.apache.tez	tez-aux-services	0.9.1.7.3.1.300-81
Tez	org.apache.tez	tez-common	0.9.1.7.3.1.300-81
Tez	org.apache.tez	tez-dag	0.9.1.7.3.1.300-81
Tez	org.apache.tez	tez-examples	0.9.1.7.3.1.300-81
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.3.1.300-81
Tez	org.apache.tez	tez-history-parser	0.9.1.7.3.1.300-81
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.3.1.300-81
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.3.1.300-81
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.3.1.300-81

Project	groupId	artifactId	version
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.3.1.300-81
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.3.1.300-81
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.3.1.300-81
Tez	org.apache.tez	tez-tests	0.9.1.7.3.1.300-81
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.3.1.300-81
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.3.1.300-81
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.3.1.300-81
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.3.1.300-81
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.3.1.300-81
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.3.1.300-81
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.3.1.300-81
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.3.1.300-81
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.3.1.300-81
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.3.1.300-81
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.3.1.300-81
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.3.1.300-81
Zeppelin	org.apache.zeppelin	zeppelin-shaded-raz	0.8.2.7.3.1.300-81
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.3.1.300-81
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.3.1.300-81
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.3.1.300-81
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.3.1.300-81
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.3.1.300-81
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.3.1.300-81
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.3.1.300-81
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.3.1.300-81
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.3.1.300-81
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.3.1.300-81
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.3.1.300-81
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.3.1.300-81
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.3.1.300-81

Cloudera Runtime 7.3.1.200-90

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	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-classification-updater	2.1.0.7.3.1.200-90

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-client-common	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-client-v1	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-client-v2	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-common	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-distro	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-docs	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-graphdb-api	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-graphdb-common	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-graphdb-janus	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-hdfs-bridge	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-index-repair-tool	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-intg	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-notification	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-plugin-classloader	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-repository	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-server-api	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	atlas-testtools	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	hbase-bridge	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	hbase-bridge-shim	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	hbase-testing-util	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	hdfs-model	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	hive-bridge	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	hive-bridge-shim	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	impala-bridge	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	impala-bridge-shim	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	impala-hook-api	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	kafka-bridge	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	kafka-bridge-shim	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	navigator-to-atlas	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	sample-app	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	sqoop-bridge	2.1.0.7.3.1.200-90
Atlas	org.apache.atlas	sqoop-bridge-shim	2.1.0.7.3.1.200-90
Avro	org.apache.avro	avro	1.11.1.7.3.1.200-90
Avro	org.apache.avro	avro-android	1.11.1.7.3.1.200-90
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.3.1.200-90
Avro	org.apache.avro	avro-compiler	1.11.1.7.3.1.200-90
Avro	org.apache.avro	avro-grpc	1.11.1.7.3.1.200-90

Project	groupId	artifactId	version
Avro	org.apache.avro	avro-ipc	1.11.1.7.3.1.200-90
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.3.1.200-90
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.3.1.200-90
Avro	org.apache.avro	avro-mapred	1.11.1.7.3.1.200-90
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.3.1.200-90
Avro	org.apache.avro	avro-perf	1.11.1.7.3.1.200-90
Avro	org.apache.avro	avro-protobuf	1.11.1.7.3.1.200-90
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.3.1.200-90
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.3.1.200-90
Avro	org.apache.avro	avro-thrift	1.11.1.7.3.1.200-90
Avro	org.apache.avro	avro-tools	1.11.1.7.3.1.200-90
Avro	org.apache.avro	trevni-avro	1.11.1.7.3.1.200-90
Avro	org.apache.avro	trevni-core	1.11.1.7.3.1.200-90
Calcite	org.apache.calcite	calcite-babel	1.25.0.7.3.1.200-90
Calcite	org.apache.calcite	calcite-core	1.25.0.7.3.1.200-90
Calcite	org.apache.calcite	calcite-druid	1.25.0.7.3.1.200-90
Calcite	org.apache.calcite	calcite-linq4j	1.25.0.7.3.1.200-90
Calcite	org.apache.calcite	calcite-server	1.25.0.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-benchmark	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.3.1.200-90

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-minicluster	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.3.1.200-90

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.3.1.200-90
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.3.1.200-90
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-client	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-common	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-examples	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-http	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-it	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-logging	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-replication	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-rest	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-server	2.4.17.7.3.1.200-90

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-shell	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.3.1.200-90
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.3.1.200-90
Hive	org.apache.hive	catalogd-unit	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-beeline	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-classification	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-cli	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-common	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-contrib	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-exec	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-iceberg-catalog	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-iceberg-handler	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-iceberg-shading	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-impala	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-it-iceberg	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-it-impala	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-it-util	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-jmh	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.3.1.200-90

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-metastore	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-parser	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-serde	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-service	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-shims	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-streaming	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-testutils	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-udf	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	kafka-handler	3.1.3000.7.3.1.200-90
Hive	org.apache.hive	patched-iceberg-api	patched-1.3.1.7.3.1.200-90-3.1.3000.
Hive	org.apache.hive	patched-iceberg-core	patched-1.3.1.7.3.1.200-90-3.1.3000.
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector-spark3_2.12	1.0.0.7.3.1.200-90
Kafka	org.apache.kafka	ci	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	connect	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	connect-api	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	connect-file	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	connect-json	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	generator	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.3.1.200-90

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.3.1.200-90
Kafka	org.apache.kafka	trogdor	3.4.1.7.3.1.200-90

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-adapter	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-admin-ui	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-applications	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-cloud-bindings	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-demo-ldap	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-demo-ldap-launcher	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-discovery-ambari	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-discovery-cm	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-docker	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-i18n	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-i18n-logging-log4j	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-i18n-logging-sl4j	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-openapi-ui	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-performance-test	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-ha	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-identity-assertion-common	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-identity-assertion-no-doas	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-jersey	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-rewrite	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-rewrite-common	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-security-authc-anon	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-security-authz-acls	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-security-authz-composite	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-security-authz-path-acls	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-security-clientcert	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-security-hadoopauth	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-security-jwt	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-security-pac4j	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-security-preauth	2.0.0.7.3.1.200-90

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-provider-security-shiro	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-provider-security-webappsec	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-release	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-server	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-server-launcher	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-server-xforwarded-filter	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-admin	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-as	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-auth	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-definitions	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-hashicorp-vault	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-hbase	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-health	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-hive	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-idbroker	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-idbroker-plugins	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-impala	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-jkg	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-knoxsso	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-knoxsout	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-knoxtoken	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-livy	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-metadata	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-nifi	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-nifi-registry	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-remoteconfig	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-rm	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-session	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-storm	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-test	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-tgs	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-vault	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-service-webhdfs	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-shell	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-shell-launcher	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-shell-release	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-shell-samples	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-spi	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-spi-common	2.0.0.7.3.1.200-90

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-test	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-test-idbroker	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-test-release-utils	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-test-utils	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-topology-hadoop-xml	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-topology-simple	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-util-common	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-util-configinjector	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-util-launcher	2.0.0.7.3.1.200-90
Knox	org.apache.knox	gateway-util-urltemplate	2.0.0.7.3.1.200-90
Knox	org.apache.knox	hadoop-examples	2.0.0.7.3.1.200-90
Knox	org.apache.knox	knox-cli-launcher	2.0.0.7.3.1.200-90
Knox	org.apache.knox	knox-homepage-ui	2.0.0.7.3.1.200-90
Knox	org.apache.knox	knox-token-generation-ui	2.0.0.7.3.1.200-90
Knox	org.apache.knox	knox-token-management-ui	2.0.0.7.3.1.200-90
Knox	org.apache.knox	knox-webshell-ui	2.0.0.7.3.1.200-90
Knox	org.apache.knox	webhdfs-kerb-test	2.0.0.7.3.1.200-90
Knox	org.apache.knox	webhdfs-test	2.0.0.7.3.1.200-90
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.3.1.200-90
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.3.1.200-90
Kudu	org.apache.kudu	kudu-client	1.17.0.7.3.1.200-90
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.3.1.200-90
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.3.1.200-90
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.3.1.200-90
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.3.1.200-90
Livy	org.apache.livy	livy-api	0.7.23000.7.3.1.200-90
Livy	org.apache.livy	livy-client-common	0.7.23000.7.3.1.200-90
Livy	org.apache.livy	livy-client-http	0.7.23000.7.3.1.200-90
Livy	org.apache.livy	livy-core_2.12	0.7.23000.7.3.1.200-90
Livy	org.apache.livy	livy-examples	0.7.23000.7.3.1.200-90
Livy	org.apache.livy	livy-integration-test	0.7.23000.7.3.1.200-90
Livy	org.apache.livy	livy-repl_2.12	0.7.23000.7.3.1.200-90
Livy	org.apache.livy	livy-rsc	0.7.23000.7.3.1.200-90
Livy	org.apache.livy	livy-scala-api_2.12	0.7.23000.7.3.1.200-90
Livy	org.apache.livy	livy-server	0.7.23000.7.3.1.200-90
Livy	org.apache.livy	livy-test-lib	0.7.23000.7.3.1.200-90
Livy	org.apache.livy	livy-thriftserver	0.7.23000.7.3.1.200-90
Livy	org.apache.livy	livy-thriftserver-session	0.7.23000.7.3.1.200-90
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.3.1.200-90

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-analyzers-openslp	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-core	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-join	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.3.1.200-90
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.3.1.200-90
Oozie	org.apache.oozie	oozie-client	5.1.0.7.3.1.200-90
Oozie	org.apache.oozie	oozie-core	5.1.0.7.3.1.200-90
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.3.1.200-90
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.3.1.200-90
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.3.1.200-90
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.3.1.200-90
Oozie	org.apache.oozie	oozie-server	5.1.0.7.3.1.200-90
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.3.1.200-90
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.3.1.200-90

Project	groupId	artifactId	version
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.3.1.200-90
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.3.1.200-90
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.3.1.200-90
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.3.1.200-90
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.3.1.200-90
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.3.1.200-90
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.3.1.200-90
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.3.1.200-90
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.3.1.200-90
ORC	org.apache.orc	orc-core	1.8.3.7.3.1.200-90
ORC	org.apache.orc	orc-examples	1.8.3.7.3.1.200-90
ORC	org.apache.orc	orc-mapreduce	1.8.3.7.3.1.200-90
ORC	org.apache.orc	orc-shims	1.8.3.7.3.1.200-90
ORC	org.apache.orc	orc-tools	1.8.3.7.3.1.200-90
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-client	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-common	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-config	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-crypto-api	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-crypto-default	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-erasurencode	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-managed-rocksdb	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-client	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-common	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.3.1.200-90

Project	groupId	artifactId	version
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-filessystem	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-filessystem-common	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-filessystem-hadoop2	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-filessystem-hadoop3	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-filessystem-shaded	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-httpsgateway	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.3.1.200-90
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.3.1.200-90
Parquet	org.apache.parquet	parquet-avro	1.12.3.7.3.1.200-90
Parquet	org.apache.parquet	parquet-cascading	1.12.3.7.3.1.200-90
Parquet	org.apache.parquet	parquet-cascading3	1.12.3.7.3.1.200-90
Parquet	org.apache.parquet	parquet-column	1.12.3.7.3.1.200-90
Parquet	org.apache.parquet	parquet-common	1.12.3.7.3.1.200-90
Parquet	org.apache.parquet	parquet-encoding	1.12.3.7.3.1.200-90
Parquet	org.apache.parquet	parquet-format-structures	1.12.3.7.3.1.200-90
Parquet	org.apache.parquet	parquet-generator	1.12.3.7.3.1.200-90
Parquet	org.apache.parquet	parquet-hadoop	1.12.3.7.3.1.200-90
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.12.3.7.3.1.200-90
Parquet	org.apache.parquet	parquet-jackson	1.12.3.7.3.1.200-90
Parquet	org.apache.parquet	parquet-pig	1.12.3.7.3.1.200-90
Parquet	org.apache.parquet	parquet-pig-bundle	1.12.3.7.3.1.200-90
Parquet	org.apache.parquet	parquet-protobuf	1.12.3.7.3.1.200-90
Parquet	org.apache.parquet	parquet-scala_2.12	1.12.3.7.3.1.200-90
Parquet	org.apache.parquet	parquet-thrift	1.12.3.7.3.1.200-90
Parquet	org.apache.parquet	parquet-tools	1.12.3.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.3.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.3.7.3.1.200-90

Project	groupId	artifactId	version
Phoenix	org.apache.phoenix	phoenix-core	5.1.3.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.1.6	5.1.3.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.2.5	5.1.3.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.3.0	5.1.3.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.0	5.1.3.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.1.3.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.0	5.1.3.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.4	5.1.3.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.3.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.3.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.3.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix5-spark3	6.0.0.7.3.1.200-90
Phoenix	org.apache.phoenix	phoenix5-spark3-shaded	6.0.0.7.3.1.200-90
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	jisql	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-gs-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.3.1.200-90

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-raz-hook-s3	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-raz-s3	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-raz-s3-lib	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-s3-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.3.1.200-90

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-trino-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-util	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	sample-client	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	sampleapp	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	shaded-raz-hook-s3	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.3.1.200-90
Ranger	org.apache.ranger	unixusersync	2.4.0.7.3.1.200-90
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.3.1.200-90
Solr	org.apache.solr	solr-analytics	8.11.2.7.3.1.200-90
Solr	org.apache.solr	solr-cell	8.11.2.7.3.1.200-90
Solr	org.apache.solr	solr-core	8.11.2.7.3.1.200-90
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.3.1.200-90
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.3.1.200-90
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.3.1.200-90
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.3.1.200-90
Solr	org.apache.solr	solr-langid	8.11.2.7.3.1.200-90
Solr	org.apache.solr	solr-ltr	8.11.2.7.3.1.200-90
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.3.1.200-90
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.3.1.200-90
Solr	org.apache.solr	solr-security-util	8.11.2.7.3.1.200-90
Solr	org.apache.solr	solr-solrj	8.11.2.7.3.1.200-90
Solr	org.apache.solr	solr-test-framework	8.11.2.7.3.1.200-90
Solr	org.apache.solr	solr-velocity	8.11.2.7.3.1.200-90
Spark	org.apache.spark	spark-avro_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-catalyst_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-common-utils_2.12	3.5.4.7.3.1.200-90

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-connect-client-jvm_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-connect-common_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-connect_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-core_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-graphx_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-hadoop-cloud_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-hive_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-kubernetes_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-kvstore_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-launcher_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-mllib-local_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-mllib_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-network-common_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-network-shuffle_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-network-yarn_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-protobuf_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-repl_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-shaded-raz	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-sketch_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-sql-api_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-sql-kafka-0-10_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-sql_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-streaming_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-tags_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-unsafe_2.12	3.5.4.7.3.1.200-90
Spark	org.apache.spark	spark-yarn_2.12	3.5.4.7.3.1.200-90
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.3.1.200-90
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.3.1.200-90
Tez	org.apache.tez	hadoop-shim	0.9.1.7.3.1.200-90
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.3.1.200-90
Tez	org.apache.tez	tez-api	0.9.1.7.3.1.200-90
Tez	org.apache.tez	tez-aux-services	0.9.1.7.3.1.200-90
Tez	org.apache.tez	tez-common	0.9.1.7.3.1.200-90
Tez	org.apache.tez	tez-dag	0.9.1.7.3.1.200-90
Tez	org.apache.tez	tez-examples	0.9.1.7.3.1.200-90
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.3.1.200-90

Project	groupId	artifactId	version
Tez	org.apache.tez	tez-history-parser	0.9.1.7.3.1.200-90
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.3.1.200-90
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.3.1.200-90
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.3.1.200-90
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.3.1.200-90
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.3.1.200-90
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.3.1.200-90
Tez	org.apache.tez	tez-tests	0.9.1.7.3.1.200-90
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.3.1.200-90
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.3.1.200-90
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.3.1.200-90
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.3.1.200-90
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.3.1.200-90
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.3.1.200-90
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.3.1.200-90
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.3.1.200-90
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.3.1.200-90
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.3.1.200-90
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.3.1.200-90
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.3.1.200-90
Zeppelin	org.apache.zeppelin	zeppelin-shaded-raz	0.8.2.7.3.1.200-90
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.3.1.200-90
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.3.1.200-90
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.3.1.200-90
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.3.1.200-90
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.3.1.200-90
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.3.1.200-90
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.3.1.200-90
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.3.1.200-90
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.3.1.200-90
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.3.1.200-90
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.3.1.200-90
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.3.1.200-90
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.3.1.200-90

Cloudera Runtime 7.3.1.100-57

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	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-classification-updater	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-client-common	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-client-v1	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-client-v2	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-common	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-distro	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-docs	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-graphdb-api	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-graphdb-common	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-graphdb-janus	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-hdfs-bridge	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-index-repair-tool	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-intg	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-notification	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-plugin-classloader	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-repository	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-server-api	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	atlas-testtools	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	hbase-bridge	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	hbase-bridge-shim	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	hbase-testing-util	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	hdfs-model	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	hive-bridge	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	hive-bridge-shim	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	impala-bridge	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	impala-bridge-shim	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	impala-hook-api	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	kafka-bridge	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	kafka-bridge-shim	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	navigator-to-atlas	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	sample-app	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	sqoop-bridge	2.1.0.7.3.1.100-57
Atlas	org.apache.atlas	sqoop-bridge-shim	2.1.0.7.3.1.100-57
Avro	org.apache.avro	avro	1.11.1.7.3.1.100-57

Project	groupId	artifactId	version
Avro	org.apache.avro	avro-android	1.11.1.7.3.1.100-57
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.3.1.100-57
Avro	org.apache.avro	avro-compiler	1.11.1.7.3.1.100-57
Avro	org.apache.avro	avro-grpc	1.11.1.7.3.1.100-57
Avro	org.apache.avro	avro-ipc	1.11.1.7.3.1.100-57
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.3.1.100-57
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.3.1.100-57
Avro	org.apache.avro	avro-mapred	1.11.1.7.3.1.100-57
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.3.1.100-57
Avro	org.apache.avro	avro-perf	1.11.1.7.3.1.100-57
Avro	org.apache.avro	avro-protobuf	1.11.1.7.3.1.100-57
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.3.1.100-57
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.3.1.100-57
Avro	org.apache.avro	avro-thrift	1.11.1.7.3.1.100-57
Avro	org.apache.avro	avro-tools	1.11.1.7.3.1.100-57
Avro	org.apache.avro	trevni-avro	1.11.1.7.3.1.100-57
Avro	org.apache.avro	trevni-core	1.11.1.7.3.1.100-57
Calcite	org.apache.calcite	calcite-babel	1.25.0.7.3.1.100-57
Calcite	org.apache.calcite	calcite-core	1.25.0.7.3.1.100-57
Calcite	org.apache.calcite	calcite-druid	1.25.0.7.3.1.100-57
Calcite	org.apache.calcite	calcite-linq4j	1.25.0.7.3.1.100-57
Calcite	org.apache.calcite	calcite-server	1.25.0.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-benchmark	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.3.1.100-57

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-miniclusterc	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.3.1.100-57

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.3.1.100-57
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.3.1.100-57
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-client	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-common	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-examples	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-http	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-it	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-logging	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-replication	2.4.17.7.3.1.100-57

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-rest	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-server	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-shell	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.3.1.100-57
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.3.1.100-57
Hive	org.apache.hive	catalogd-unit	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-beeline	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-classification	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-cli	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-common	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-contrib	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-exec	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-iceberg-catalog	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-iceberg-handler	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-iceberg-shading	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-impala	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-it-iceberg	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-it-impala	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-it-util	3.1.3000.7.3.1.100-57

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-jmh	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-metastore	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-parser	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-serde	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-service	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-shims	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-streaming	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-testutils	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-udf	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	kafka-handler	3.1.3000.7.3.1.100-57
Hive	org.apache.hive	patched-iceberg-api	patched-1.3.1.7.3.1.100-57-3.1.3000.
Hive	org.apache.hive	patched-iceberg-core	patched-1.3.1.7.3.1.100-57-3.1.3000.
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector-spark3_2.12	1.0.0.7.3.1.100-57
Kafka	org.apache.kafka	ci	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	connect	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	connect-api	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	connect-file	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	connect-json	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.3.1.100-57

Project	groupId	artifactId	version
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	generator	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.3.1.100-57
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.3.1.100-57

Project	groupId	artifactId	version
Kafka	org.apache.kafka	trogdor	3.4.1.7.3.1.100-57
Knox	org.apache.knox	gateway-adapter	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-admin-ui	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-applications	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-cloud-bindings	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-demo-ldap	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-demo-ldap-launcher	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-discovery-ambari	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-discovery-cm	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-docker	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-i18n	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-i18n-logging-log4j	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-i18n-logging-sl4j	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-openapi-ui	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-performance-test	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-ha	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-identity-assertion-common	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-identity-assertion-no-doas	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-jersey	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-rewrite	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-rewrite-common	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-security-authc-anon	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-security-authz-acls	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-security-authz-composite	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-security-authz-path-acls	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-security-clientcert	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-security-hadoopauth	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-security-jwt	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-security-pac4j	2.0.0.7.3.1.100-57

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-provider-security-preauth	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-security-shiro	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-provider-security-webappsec	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-release	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-server	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-server-launcher	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-server-xforwarded-filter	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-admin	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-as	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-auth	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-definitions	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-hashicorp-vault	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-hbase	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-health	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-hive	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-idbroker	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-idbroker-plugins	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-impala	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-jkg	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-knoxsso	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-knoxssout	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-knoxtoken	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-livy	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-metadata	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-nifi	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-nifi-registry	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-remoteconfig	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-rm	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-session	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-storm	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-test	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-tgs	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-vault	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-service-webhdfs	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-shell	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-shell-launcher	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-shell-release	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-shell-samples	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-spi	2.0.0.7.3.1.100-57

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-spi-common	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-test	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-test-idbroker	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-test-release-utils	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-test-utils	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-topology-hadoop-xml	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-topology-simple	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-util-common	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-util-configinjector	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-util-launcher	2.0.0.7.3.1.100-57
Knox	org.apache.knox	gateway-util-urltemplate	2.0.0.7.3.1.100-57
Knox	org.apache.knox	hadoop-examples	2.0.0.7.3.1.100-57
Knox	org.apache.knox	knox-cli-launcher	2.0.0.7.3.1.100-57
Knox	org.apache.knox	knox-homepage-ui	2.0.0.7.3.1.100-57
Knox	org.apache.knox	knox-token-generation-ui	2.0.0.7.3.1.100-57
Knox	org.apache.knox	knox-token-management-ui	2.0.0.7.3.1.100-57
Knox	org.apache.knox	knox-webshell-ui	2.0.0.7.3.1.100-57
Knox	org.apache.knox	webhdfs-kerb-test	2.0.0.7.3.1.100-57
Knox	org.apache.knox	webhdfs-test	2.0.0.7.3.1.100-57
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.3.1.100-57
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.3.1.100-57
Kudu	org.apache.kudu	kudu-client	1.17.0.7.3.1.100-57
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.3.1.100-57
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.3.1.100-57
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.3.1.100-57
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.3.1.100-57
Livy	org.apache.livy	livy-api	0.7.23000.7.3.1.100-57
Livy	org.apache.livy	livy-client-common	0.7.23000.7.3.1.100-57
Livy	org.apache.livy	livy-client-http	0.7.23000.7.3.1.100-57
Livy	org.apache.livy	livy-core_2.12	0.7.23000.7.3.1.100-57
Livy	org.apache.livy	livy-examples	0.7.23000.7.3.1.100-57
Livy	org.apache.livy	livy-integration-test	0.7.23000.7.3.1.100-57
Livy	org.apache.livy	livy-repl_2.12	0.7.23000.7.3.1.100-57
Livy	org.apache.livy	livy-rsc	0.7.23000.7.3.1.100-57
Livy	org.apache.livy	livy-scala-api_2.12	0.7.23000.7.3.1.100-57
Livy	org.apache.livy	livy-server	0.7.23000.7.3.1.100-57
Livy	org.apache.livy	livy-test-lib	0.7.23000.7.3.1.100-57
Livy	org.apache.livy	livy-thriftserver	0.7.23000.7.3.1.100-57
Livy	org.apache.livy	livy-thriftserver-session	0.7.23000.7.3.1.100-57

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-analyzers-openslp	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-core	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-join	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.3.1.100-57
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.3.1.100-57
Oozie	org.apache.oozie	oozie-client	5.1.0.7.3.1.100-57
Oozie	org.apache.oozie	oozie-core	5.1.0.7.3.1.100-57
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.3.1.100-57
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.3.1.100-57
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.3.1.100-57
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.3.1.100-57
Oozie	org.apache.oozie	oozie-server	5.1.0.7.3.1.100-57
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.3.1.100-57

Project	groupId	artifactId	version
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.3.1.100-57
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.3.1.100-57
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.3.1.100-57
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.3.1.100-57
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.3.1.100-57
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.3.1.100-57
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.3.1.100-57
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.3.1.100-57
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.3.1.100-57
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.3.1.100-57
ORC	org.apache.orc	orc-core	1.8.3.7.3.1.100-57
ORC	org.apache.orc	orc-examples	1.8.3.7.3.1.100-57
ORC	org.apache.orc	orc-mapreduce	1.8.3.7.3.1.100-57
ORC	org.apache.orc	orc-shims	1.8.3.7.3.1.100-57
ORC	org.apache.orc	orc-tools	1.8.3.7.3.1.100-57
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-client	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-common	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-config	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-crypto-api	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-crypto-default	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-managed-rocksdb	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-client	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-common	1.4.0.7.3.1.100-57

Project	groupId	artifactId	version
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-filessystem	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-filessystem-common	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-filessystem-hadoop2	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-filessystem-hadoop3	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-filessystem-shaded	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-httpsgateway	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.3.1.100-57
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.3.1.100-57
Parquet	org.apache.parquet	parquet-avro	1.12.3.7.3.1.100-57
Parquet	org.apache.parquet	parquet-cascading	1.12.3.7.3.1.100-57
Parquet	org.apache.parquet	parquet-cascading3	1.12.3.7.3.1.100-57
Parquet	org.apache.parquet	parquet-column	1.12.3.7.3.1.100-57
Parquet	org.apache.parquet	parquet-common	1.12.3.7.3.1.100-57
Parquet	org.apache.parquet	parquet-encoding	1.12.3.7.3.1.100-57
Parquet	org.apache.parquet	parquet-format-structures	1.12.3.7.3.1.100-57
Parquet	org.apache.parquet	parquet-generator	1.12.3.7.3.1.100-57
Parquet	org.apache.parquet	parquet-hadoop	1.12.3.7.3.1.100-57
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.12.3.7.3.1.100-57
Parquet	org.apache.parquet	parquet-jackson	1.12.3.7.3.1.100-57
Parquet	org.apache.parquet	parquet-pig	1.12.3.7.3.1.100-57
Parquet	org.apache.parquet	parquet-pig-bundle	1.12.3.7.3.1.100-57
Parquet	org.apache.parquet	parquet-protobuf	1.12.3.7.3.1.100-57
Parquet	org.apache.parquet	parquet-scala_2.12	1.12.3.7.3.1.100-57
Parquet	org.apache.parquet	parquet-thrift	1.12.3.7.3.1.100-57
Parquet	org.apache.parquet	parquet-tools	1.12.3.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.3.7.3.1.100-57

Project	groupId	artifactId	version
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.3.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix-core	5.1.3.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.1.6	5.1.3.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.2.5	5.1.3.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.3.0	5.1.3.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.0	5.1.3.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.1.3.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.0	5.1.3.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.4	5.1.3.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.3.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.3.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.3.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix5-spark3	6.0.0.7.3.1.100-57
Phoenix	org.apache.phoenix	phoenix5-spark3-shaded	6.0.0.7.3.1.100-57
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	jisql	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-gs-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.3.1.100-57

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-raz-hook-s3	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-raz-s3	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-raz-s3-lib	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-s3-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.3.1.100-57

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-trino-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-util	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	sample-client	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	sampleapp	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	shaded-raz-hook-s3	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.3.1.100-57
Ranger	org.apache.ranger	unixusersync	2.4.0.7.3.1.100-57
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.3.1.100-57
Solr	org.apache.solr	solr-analytics	8.11.2.7.3.1.100-57
Solr	org.apache.solr	solr-cell	8.11.2.7.3.1.100-57
Solr	org.apache.solr	solr-core	8.11.2.7.3.1.100-57
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.3.1.100-57
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.3.1.100-57
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.3.1.100-57
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.3.1.100-57
Solr	org.apache.solr	solr-langid	8.11.2.7.3.1.100-57
Solr	org.apache.solr	solr-ltr	8.11.2.7.3.1.100-57
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.3.1.100-57
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.3.1.100-57
Solr	org.apache.solr	solr-security-util	8.11.2.7.3.1.100-57
Solr	org.apache.solr	solr-solrj	8.11.2.7.3.1.100-57
Solr	org.apache.solr	solr-test-framework	8.11.2.7.3.1.100-57
Solr	org.apache.solr	solr-velocity	8.11.2.7.3.1.100-57
Spark	org.apache.spark	spark-avro_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-catalyst_2.12	3.4.1.7.3.1.100-57

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-connect-client-jvm_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-connect-common_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-connect_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-core_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-graphx_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-hadoop-cloud_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-hive_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-kubernetes_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-kvstore_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-launcher_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-mllib-local_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-mllib_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-network-common_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-network-shuffle_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-network-yarn_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-protobuf_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-repl_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-shaded-raz	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-sketch_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-sql-kafka-0-10_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-sql_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-streaming_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-tags_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-unsafe_2.12	3.4.1.7.3.1.100-57
Spark	org.apache.spark	spark-yarn_2.12	3.4.1.7.3.1.100-57
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.3.1.100-57
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.3.1.100-57
Tez	org.apache.tez	hadoop-shim	0.9.1.7.3.1.100-57
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.3.1.100-57
Tez	org.apache.tez	tez-api	0.9.1.7.3.1.100-57
Tez	org.apache.tez	tez-aux-services	0.9.1.7.3.1.100-57
Tez	org.apache.tez	tez-common	0.9.1.7.3.1.100-57
Tez	org.apache.tez	tez-dag	0.9.1.7.3.1.100-57
Tez	org.apache.tez	tez-examples	0.9.1.7.3.1.100-57
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.3.1.100-57
Tez	org.apache.tez	tez-history-parser	0.9.1.7.3.1.100-57

Project	groupId	artifactId	version
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.3.1.100-57
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.3.1.100-57
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.3.1.100-57
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.3.1.100-57
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.3.1.100-57
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.3.1.100-57
Tez	org.apache.tez	tez-tests	0.9.1.7.3.1.100-57
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.3.1.100-57
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.3.1.100-57
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.3.1.100-57
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.3.1.100-57
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.3.1.100-57
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.3.1.100-57
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.3.1.100-57
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.3.1.100-57
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.3.1.100-57
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.3.1.100-57
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.3.1.100-57
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.3.1.100-57
Zeppelin	org.apache.zeppelin	zeppelin-shaded-raz	0.8.2.7.3.1.100-57
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.3.1.100-57
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.3.1.100-57
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.3.1.100-57
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.3.1.100-57
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.3.1.100-57
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.3.1.100-57
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.3.1.100-57
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.3.1.100-57
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.3.1.100-57
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.3.1.100-57
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.3.1.100-57
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.3.1.100-57
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.3.1.100-57

Cloudera 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

Project	groupId	artifactId	version
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
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

Project	groupId	artifactId	version
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
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

Project	groupId	artifactId	version
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
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

Project	groupId	artifactId	version
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
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

Project	groupId	artifactId	version
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
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

Project	groupId	artifactId	version
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
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

Project	groupId	artifactId	version
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
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

Project	groupId	artifactId	version
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
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

Project	groupId	artifactId	version
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
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-knoxsout	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

Project	groupId	artifactId	version
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
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

Project	groupId	artifactId	version
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-opennlp	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
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

Project	groupId	artifactId	version
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
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-erasurencode	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

Project	groupId	artifactId	version
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
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

Project	groupId	artifactId	version
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
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

Project	groupId	artifactId	version
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
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

Project	groupId	artifactId	version
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
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

Project	groupId	artifactId	version
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
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

Project	groupId	artifactId	version
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
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

Release Matrix

You must be familiar with the versions of all the service packs and cumulative hotfixes for Cloudera Runtime 7.3.1.

The Release build number has two parts, the [**Runtime release version number**] and the [**Runtime Build number**]. For example, in 7.3.1.100-57, 7.3.1.100 is the release version number and 57 is the build number.

Runtime Release Build Number	Cumulative Hotfix/Service Pack	Release Date	Download URL
7.3.1.500-182	7.3.1.500 SP3	Nov 5, 2025	https://archive.cloudera.com/p/cdh7/7.3.1.500/
7.3.1.400-100	7.3.1.400 SP2	June 27, 2025	https://archive.cloudera.com/p/cdh7/7.3.1.400/
7.3.1.300-81	7.3.1.300 SP1 CHF 1	June 02, 2025	https://archive.cloudera.com/p/cdh7/7.3.1.300/
7.3.1.200-90	7.3.1.200 SP1	April 24, 2025	https://archive.cloudera.com/p/cdh7/7.3.1.200/
7.3.1.100-57	7.3.1.100 CHF 1	March 18, 2025	https://archive.cloudera.com/p/cdh7/7.3.1.100/
7.3.1.0-197	7.3.1	December 11, 2024	https://archive.cloudera.com/p/cdh7/7.3.1/

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, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

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

7.3.1.500

Atlas stores all timestamps in UTC, but the UI or API would interpret TODAY or YESTERDAY based on the local server time zone. For instance, if the server is in a different time zone from the user, TODAY may refer to different times than expected, causing mismatched results. After the update, search results remain accurate despite the server's timezone as the date and time conversions while fetching the results use UTC time zone.

CDPD-84502: Advanced Search is not working properly in the Atlas UI

7.3.1.500

Fixed an issue in the Atlas UI where Advanced Search did not trigger API calls for each relation. The UI now ensures that API calls are made for all relevant relations, providing accurate and complete search results.

This update also incorporates improvements from CDPD-65619, which enhanced the reliability and consistency of Advanced Search operations in the Atlas UI.

CDPD-79099: Unable to import hive tables using import script in Public Cloud

7.3.1.500

The validation of import .zip files is updated to prevent the org.apache.atlas.AtlasServiceException error if the .zip file contains Atlas shell entities missing mandatory attributes such as a name. The validation of the mandatory attributes is removed. The update ensures that the imported Hive table entities have the correct name and are marked as complete.

Cloudera Runtime 7.3.1.400 SP2

CDPD-82054: UI: when server response date fields as '0', UI shows as current time

7.3.1.400

Previously, if an API response contained an invalid date value (such as 0) intended for display on the user interface, the current system date was shown. Currently, instead of displaying the system date when an invalid date value is encountered in the API response, the user interface displays NA. This issue specifically affects the **Entity Detail** page, where the create time and modified time are displayed.

Apache JIRA: [ATLAS-5015](#)

CDPD-34895: Investigate which HTTP security headers Atlas should have in API responses

7.3.1.400

A static header and a custom header is provided which can be added through atlas-application.properties file. This enables adding custom security headers.

Use the following prefix keyword to customize the header:

In atlas-application.properties instead of using the prefix atlas.rest.headers use atlas.header.

Example:

```
- atlas.headers.content-security-policy=default-src 'self' // content-security-policy
```

CDPD-74374: Patch to replace the long strings set in spark_process attributes

7.3.1.400

As the atlas.process.spark.attributes.update.patch is set to TRUE by default, the spark_process entity attributes details and sparkPlanDescription no longer cause out-of-memory issues as they are no longer contain a large amount of data.

OPSAPS-74147: ZDU | Upgrade from 7.1.7.3000 to 7.3.1.400 no longer fails

7.3.1.400

The ZDU rolling upgrade for Atlas failed from Cloudera Runtime 7.1.7 SP3 to 7.3.1.x version, as the Zero Downtime Upgrade feature was not available for 7.1.7 SP3. This issue is fixed. The Role State for Atlas is now checked and the Update Solr Configs for Atlas command no longer fails.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

CDPD-80922: Without a permission for one glossary, the /glossary call throws exception and it does not list the remaining glossaries

7.3.1.300

The getGlossaries method in GlossaryService is updated to ensure that the full paginated list is retrieved, even if some glossaries are skipped. This method includes the following improvements:

- Handling skipped glossaries: If some entities fail to load, it fetches additional entities until the requested limit is met.
- Efficient pagination: Keeps fetching until it gets the required number of valid glossaries.
- Preventing infinite loops: method stops when either the required number of glossaries is retrieved or there are no more to fetch.

CDPD-70450: Impala SQL queries that include the WITH clause should populate lineage in Atlas

7.3.1.300

Previously, only Impala SQL queries that don't use the "WITH" clause could be shown with their lineage in Atlas, but queries that do use the "WITH" clause could not be shown with lineage in Atlas.

Currently, Impala SQL queries using the "WITH" clause are supported.

CDPD-80160: Use the centralized commons-lang3

7.3.1.300

A centralized commons-lang3 package is used to prevent build issues and a security vulnerability.

Cloudera Runtime 7.3.1.200 SP1

CDPD-76789: Creating tag with name description throws java.lang.ClassCastException

7.3.1.200

A validation is added to avoid creating classification with reserved names such as "name", "description", "owner", "version", "serviceType" and "options".

Apache Jira: [ATLAS-4956](#)

CDPD-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails

7.3.1.200

Running an Import-hive action on a RAZ-enabled Cloudera Data Engineering cluster no longer fails with the error NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl.

CDPD-65806: Not all Iceberg table relationships are visible after upgrade from Cloudera Public Cloud 7.2.17 to Cloudera Public Cloud 7.2.18

7.3.1.200

Iceberg table relationships like Iceberg table DB, Iceberg table SD, Iceberg DDL queries were added for Iceberg table entity. Old relationships such as Hive table DB, Hive table SD, Hive DDL queries were removed.

Cloudera Runtime 7.3.1.100 CHF 1

CDPD-77767: Migration status is not updated on file while migration is in progress

7.3.1.100

Previously, whenever a migration was started on any environment and tried to update the vertex with details, the status update would fail with a null pointer exception. This issue is fixed by using the file hash code instead of file name causing the null pointer exception. Now, whenever a migration restarts, it starts from last point where it was stopped or interrupted.

Apache Jira: [ATLAS-4907](#)

CDPD-69910: Nullpointer exception while deleting business metadata

7.3.1.100

The migration status is stored correctly and reused to restart migration from the point where it failed earlier.

Apache Jira: [ATLAS-4863](#)

CDPD-74140: Comment for existing Hive tables can not be modified in Atlas

7.3.1.100

The comment attribute is now updated and existing comments can be modified and the changes are reflected in the Apache Atlas UI.

CDPD-76536: Iceberg entities created in 7.2.17 Atlas-Hook (Hive, Impala, Spark) are not ingested by 7.2.18/7.3.1 Atlas Server

7.3.1.100

The incompatibility issue between Cloudera Data Services for CDP Public Cloud 7.2.17 Apache Hive and Apache Impala hooks with Atlas server for CDP Public Cloud 7.2.18 and Cloudera on cloud 7.3.1 is resolved. With this fix, Iceberg tables created in Cloudera Public Cloud 7.2.17 or older are captured properly in Apache Atlas UI for Cloudera Public Cloud 7.2.18 and Cloudera on cloud 7.3.1.

CDPD-65619: Iceberg tables do not show under hive_db entity when created

7.3.1.100

The UI now displays Iceberg tables if they are created under the same Hive database (hive_db). Iceberg table can be seen in the **Tables** tab of a Hive database (hive_db) in basic search. Both Hive table (hive_table) and Iceberg table (iceberg_table) typenames are visible.

CDPD-71411: Atlas Hbase import fails with jackson-databind exception

7.3.1.100

The jackson-annotations is updated to match jackson-databind (2.15.0) to resolve the NoSuchFieldError error during an HBase import caused by version mismatch (older 2.12.7 on classpath).

Cloudera Runtime 7.3.1**CDPD-69962: fetchType as "incremental" does full export instead of "CONNECTED"**

7.3.1

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

7.3.1

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 Cloudera Base on premises

7.3.1

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

7.3.1

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 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

Review the list of Avro issues that are resolved in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.400 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.200 SP1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.100 CHF 1

CDPD-47852: Removed the old CDH versions and parent

7.3.1.100

Removed old CDH repository reference for artefacts to support build with newer dependency versions.

CDPD-45628: Upgraded Apache Maven to 3.8.6 due to CVE-2021-26291

7.3.1.100

Removed the Maven prerequisites of version 2.2.1 and upgrade the maven-core to 3.8.6 to fix CVE-2021-26291. Also, upgraded plexus-utils version to 3.5.0 and Apache file-management version to 3.0.0 to support the upgrade.

CDPD-75089: Restrict trusted packages in ReflectData and SpecificData

7.3.1.100

Schema parsing in Java SDK of Apache Avro had an issue that allowed malicious actors to execute arbitrary code when reading Avro data. This issue is now resolved by restricting trusted packages in ReflectData and SpecificData.

Apache Jira:[AVRO-3985](#)

Cloudera Runtime 7.3.1

There are no fixed issues in this release.

Fixed issues in Cloud Connectors

Review the list of Cloud Connectors issues that are resolved in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.400 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.200 SP1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.100 CHF 1**CDPD-76378 Uploading files to S3 takes longer than expected**

7.3.1.100

The issue regarding the delayed upload of files in S3 is fixed.

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

Review the list of Cruise Control issues that are resolved in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.400 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1**CDPD-75633: Cruise Control might fail on startup or create metrics topic with default configuration if Kafka brokers lag**

7.3.1.300

When first starting up Cruise Control and Kafka brokers, one of the brokers might lag behind, or the cluster might not be ready to create the metrics topic correctly, which previously caused Cruise Control to fail on startup. This no longer occurs because topic creation is now checked and retried if necessary during startup.

Cloudera Runtime 7.3.1.200 SP1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.100 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1

There are no fixed issues in this release.

Fixed Issues in HDFS

Review the list of HDFS issues that are resolved in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

CDPD-84568: Backport upstream fixes for missing blocks issue

7.3.1.500

Backported the following upstream fixes:

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

Cloudera Runtime 7.3.1.400 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

CDPD-77765: Hadoop - Upgrade kafka-clients to 3.7.2/3.8.1+ due to CVE-2024-56128

7.3.1.300

Upgrade kafka client to 3.9.0 to fix CVE-2024-31141.

Apache Jira: [HADOOP-19456](#)

Cloudera Runtime 7.3.1.200 SP1

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

7.3.1.200

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

Apache Jira: [HADOOP-16299](#)

CDPD-75403: Temporary File Local Information Disclosure

7.3.1.200

The issue is fixed and this fix enhances access control for the current directory of the RunJar so that only current user can have the access.

Apache Jira: [HADOOP-19031](#)

CDPD-69173: FsImage might have inaccessible nodes

7.3.1.200

If a FsImage is corrupted, it may have inaccessible nodes. The FsImageValidation tool currently is able to identify the inaccessible nodes when validating the INodeMap. This fix updates the FsImageValidation to remove the inaccessible nodes.

Apache Jira: [HDFS-17380](#)

Cloudera Runtime 7.3.1.100 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1

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

7.3.1

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: Cloudera Operational Database cluster creation is failing on INT and displays the Failed to create HDFS directory /tmp error.

7.3.1

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.

7.3.1

This issue is now resolved.

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

7.3.1

This issue is now resolved by fixinf service.sdl.

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

7.3.1

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, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.400 SP2

CDPD-84435: The upgrade operation fails with a message “Failed to decommission RegionServer.”

7.3.1.400

This issue is fixed. Now, the HBase MASTER aborts when it detects a `WALSyncTimeoutException` while making edits to the MasterRegion.

Apache Jira: [HBASE-28803](#)

CDPD-83544 : Potential performance degradation may occur when utilizing persistent cache, causing a restart involving full cache recovery.

7.3.1.400

This issue is fixed.

Apache Jira: [HBASE-29326](#)

CDPD-81524: Add configurable throttling of region moves in CacheAwareLoadBalancer

7.3.1.400

This fix introduces region moving throttling for LoadBalancer implementations. The throttling time is configurable by the `hbase.master.balancer.move.throttlingMillis` property, with a default value of 60000 milliseconds.

In this change, the only balancer implementation applying throttling is the `CacheAwareLoadBalancer`. All other balancers just inherit the noop default provided within the `LoadBalancer` interface.

The `CacheAwareLoadBalancer` throttling implementation performs throttling only for regions moving to the target server with a region cached ratio below the threshold configurable by `hbase.master.balancer.stochastic.throttling.cacheRatio` (80% by default).

Apache Jira: [HBASE-29168](#)

CDPD-81524: The ``ENCODED_DATA`` block type is not being considered within ``BucketCache.notifyFileCachingComplete``

7.3.1.400

This fix addresses a defect in `BucketCache.notifyFileCachingComplete`, wherein only blocks of the `DATA` type were registered. When an encoding such as `FASTDIFF` was employed, the data block type became `ENCODED_DATA`, preventing it from being accounted for in the internal cache metrics. This oversight subsequently affects the cache-aware balancer after cache recovery following a crash or restart (with persistent cache enabled), as the region percentage cache is not accurately calculated due to this flaw.

Apache Jira: [HBASE-29243](#)

CDPD-81524: Enable `BlockCache` implementations to define dynamic properties

7.3.1.400

This resolution introduces dynamic configurability for the following properties related to free space management and block prioritization:

- `hbase.bucketcache.acceptfactor`
- `hbase.bucketcache.minfactor`
- `hbase.bucketcache.extrafreefactor`
- `hbase.bucketcache.single.factor`
- `hbase.bucketcache.multi.factor`
- `hbase.bucketcache.multi.factor`
- `hbase.bucketcache.memory.factor`
- `hbase.bucketcache.queue.addition.waittime`
- `hbase.bucketcache.persist.intervalinmillis`
- `hbase.bucketcache.persistence.chunksize`

Apache Jira: [HBASE-29249](#)

CDPD-81524: Display hit ratio metrics by configurable, granular periods

7.3.1.400

This change introduces two additional properties:

- `hbase.blockcache.stats.periods` which allows defining a multiple window period;
- `hbase.blockcache.stats.period.minute` which defines the length of each of these periods (in minutes);

If `hbase.blockcache.stats.periods` is defined and is greater than one, it creates a scheduled executor that rolls the metrics calculation at `hbase.blockcache.stats.period.minute` rate. This property calculates the hit ratio for each of the last periods (as defined by `hbase.blockcache.stats.periods`), accounting for only the hits and requests that occurred during the interval of the given period (as defined by `hbase.blockcache.stats.period.minute`).

Apache Jira: [HBASE-29276](#)

CDPD-81524: Avoid adding new blocks during prefetch if usage is greater than the accept factor

7.3.1.400

Previously, when cache prefetch was enabled and cache usage reached the configured acceptance factor, it resulted in a cycle of frequent mass block evictions until the prefetch thread completed reading the entire file. This process proved to be both costly and inefficient. An initial attempt to mitigate this issue was proposed in [HBASE-28176](#); however, that solution only interrupted the prefetch thread after it had already attempted to cache the current block being read, which could still trigger a mass eviction.

To completely avert evictions triggered solely by the prefetch, this modification evaluates the impact of incorporating the current block into the cache before attempting to write it into the cache. This verification is exclusively executed when caching from prefetch threads; standard client reads and HFile writes persist in their attempt to cache the associated block.

Apache Jira: [HBASE-29288](#)

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.200 SP1**CDPD-77399: HBase fails to register the servlet metrics and throws ClassNotFoundException: org.apache.hadoop.metrics.MetricsServlet**

7.3.1.200

This issue is fixed now. HBase does not warn about the Hadoop 2-based metric servlet class on a Hadoop 3 deployment.

Apache Jira: [HBASE-28315](#)

Cloudera Runtime 7.3.1.100 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1**CDPD-67520: JWT authentication expects [sub] claim in the payload**

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

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

7.3.1

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

7.3.1

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

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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)

7.3.1

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

7.3.1

This issue is fixed.

Apache Jira: [HBASE-28500](#)

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

7.3.1

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

7.3.1

This issue is fixed.

Apache Jira: [HBASE-28501](#)

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

7.3.1

This issue is fixed.

Apache Jira: [HBASE-28626](#)

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

7.3.1

This issue is fixed.

Apache Jira: [HBASE-28523](#)

CDPD-68900: HBase properties need to be dynamically configured

7.3.1

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

7.3.1

This issue is fixed.

Apache Jira: [HBASE-28458](#)

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

7.3.1

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

7.3.1

This issue is fixed.

Apache Jira: [HBASE-28211](#)

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

7.3.1

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

7.3.1

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

7.3.1

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, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

CDPD-83511/CDPD-81079: Metastore did not enforce maximum thrift message size

7.3.1.500

The Metastore server always used the default 100 MB Thrift message size, even if a higher limit was set on the client. Large client requests caused silent connection drops and unclear exceptions.

The issue was addressed by applying the configured maximum Thrift message size on the server, ensuring consistent behavior and avoiding unexpected disconnections.

Apache Jira: [HIVE-28824](#)

CDPD-85632: Wrong results when CASE expressions has function calls referencing *CHAR* type expressions or columns

7.3.1.500

Queries using CASE expressions with nested function calls (such as UPPER()) having *CHAR* type expressions or columns as parameters returned incorrect results due to type mismatches and whitespace handling during execution. Example:

```
case upper(col1) when 'A' then 'OK' else 'N/A' end Where col1 is  
a CHAR type column.
```

The issue was addressed by ensuring consistent type casting during query planning so that CHAR values are compared correctly after conversion.

Apache Jira: [HIVE-28792](#)

CDPD-85631: Jetty version

7.3.1.500

Jetty version was upgraded to 9.4.57.v20241219 and mina-core pkg was excluded to fix CVEs.

CDPD-84566: Stable Oozie Hive actions for DML operations

7.3.1.500

Oozie Hive actions for DML operations previously failed because a Tez session was unnecessarily started during Hive CLI startup. The Oozie action would complete and delete temporary directories before the Tez session could finish, causing the session to fail with an exception:

```
FileNotFoundException
```

.

This issue is now resolved by adding a setting to disable the premature opening of a Tez session during Hive CLI startup. This prevents the unnecessary session from failing and ensures Oozie Hive actions for DML operations complete successfully.

Apache Jira: [HIVE-27023](#)

CDPD-84564: Improved query error reporting for analysis

7.3.1.500

Query error messages were previously not saved at the driver level, which made it difficult to analyze the reason for a simple query failure for later reference or query-tracking purposes.

This issue is now resolved by storing the query error message in a variable within the DriverContext.

Apache Jira: [HIVE-28312](#)

CPDP-60943: Data Loss during compaction

7.3.1.500

Data loss no longer occurs during compaction when Apache Ranger policies for masking or row filtering are enabled and compaction users are included in the policies. Compaction queries are now automatically excluded from ALL ranger policies.

Apache Jira: [HIVE-27643](#)

Housekeeping task to clean up local folders

7.3.1.500

When a Hive LLAP daemon crashes, it can leave behind unnecessary files in the LLAP local directories, which can lead to disk overflow and cause Hive queries to fail with an invalid disk error exception.

A new housekeeping task has been added to clean up these local folders on Hive Virtual Warehouse startup and periodically after startup. This resolves the disk overflow issue. The following properties are introduced to manage this cleanup process:

- `hive.llap.local.dir.cleaner.cleanup.interval`: Specifies the time interval for the LocalDirCleaner service to check for stale files. Default value: 2 hours.
- `hive.llap.local.dir.cleaner.file.modify.time.threshold`: Specifies the threshold time. Any file older than this time is deleted. Default value: 24 hours.

Incremental rebuild accuracy for materialized views

7.3.1.500

Optimized the accuracy of representing incremental rebuild availability for materialized views containing aggregate functions in their definition query.

Apache Jira: [HIVE-28006](#)

Fix for partition metadata fetch issue during DST shift

7.3.1.500

Fetching partition metadata from Hive Metastore using Java Data Objects (JDO) by timestamp now does not provide incorrect results in daylight saving time (DST) shift in partition pruning.

Apache Jira: [HIVE-27775](#)

Materialized view rebuild issue with delete operations

7.3.1.500

Addressed data correctness issues when rebuilding a materialized view incrementally in scenarios involving delete operations in one of its source tables. This is evident when the materialized view definition query includes joins and aggregates.

Apache Jira: [HIVE-27924](#)

CPDP-66779: Partitioned Iceberg table not getting loaded with insert select query from Hive

7.3.1.500

If a partitioned table is created in Iceberg and a user attempts to insert data from another table using an `insert into ... select` query, an error occurs.

This issue is now fixed.

CPDP-74640: Improved query consistency and data writing for Beeline and Hive queries

7.3.1.500

In concurrent workflows using Beeline, queries occasionally returned incorrect results due to non-thread-safe file handling, especially when `hive.query.result.cached.enabled` was disabled.

Additionally, INSERT OVERWRITE DIRECTORY operations failed to write data correctly to specified directories when query result caching was enabled.

The issue was addressed by implementing thread-safe file handling in HiveSequenceFileInputFormat and adjusting cache handling for directory overwrite queries, ensuring reliable query results and consistent data writes in concurrent workflows.

Apache Jira: [HIVE-28530](#), [HIVE-21386](#), [HIVE-25907](#)

CDPD-72985: Compatibility issue in HMS thrift struct for Hive column stats

7.3.1.500

Hive4 introduced a new required "engine" field to differentiate the stats generated by different engines. This broke compatibility with clients using Hive 3 or other engines using customized thrift API, such as TrinoDB.

The "engine" field in the Hive Metastore (HMS) API was made optional, which restores compatibility with Hive 3 clients for column stats retrieval.

Apache Jira: [HIVE-27984](#)

CDPD-71484: Improve LLAP performance by reusing FileSystem objects across tasks

7.3.1.500

Frequent closure of FileSystem objects disabled Hadoop's FileSystem cache, reducing LLAP efficiency.

Adjusted FileSystem handling to close objects once per query and daemon rather than per task, enhancing reuse and maintaining cache functionality.

Apache Jira: [HIVE-27884](#)

CDPD-70956: Queries over JDBC tables fail due to column types mismatch

7.3.1.500

Queries over JDBC tables fail at runtime when there is a mismatch between the Hive type and the database type for some columns and the Cost-Based Optimizer (CBO) is not used.

Apache Jira: [HIVE-28285](#)

DWX-17619: HPL/SQL built-in function unexpected output

7.3.1.500

Certain HPL/SQL built-in functions, such as lower and trim, were not functioning correctly when used in INSERT statements. This issue occurred after a previous fix that removed UDFs required for HPL/SQL's local and offline modes.

The issue was resolved by re-adding the necessary UDFs to HPL/SQL to ensure compatibility with local and offline modes. Related issues with these UDFs were also fixed to restore their functionality in INSERT and other SELECT statements.

Apache Jira: [HIVE-28143](#)

CDPD-74205: SharedWorkOptimizer leaves residual unused operator tree

7.3.1.500

The shared work optimizer left behind unused operator trees that sent dynamic partition pruning (DPP) events to non-existing table scan operators. This caused errors during query execution, such as 'No work found for tablescan TS[53]', disrupting workflows and query processing.

The issue was fixed by removing any leftover operator trees that sent dynamic partition pruning events to unknown operators during the optimization process. The fix ensures smoother query execution and prevents such errors.

Apache Jira: [HIVE-28484](#)

CDPD-73269: RexLiteral to ExprNode conversion issue with empty string

7.3.1.500

The conversion from RexLiteral to ExprNode failed when the literal was an empty string. This issue, introduced in HIVE-23892, caused the Cost-Based Optimizer (CBO) to fail for queries containing filters with empty literals.

The issue was fixed by ensuring that an empty literal in the filter still produces a valid RexNode during the conversion process. This fix prevents CBO failures for such queries.

Apache Jira: [HIVE-28431](#)

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

7.3.1.500

The absence of metastore.storage.schema.reader.impl caused Avro table import or download failures in Cloudera 7.1.7 when using the ODBC driver. The issue was addressed by ensuring that all records are correctly preserved during major compaction.

Apache Jira: [HIVE-26952](#)

CDPD-72605: Optimized partition authorization in HiveMetaStore to reduce overhead

7.3.1.500

The add_partitions() API in HiveMetaStore was authorizing both new and existing partitions, leading to unnecessary processing and increased load on the authorization service.

The issue was addressed by modifying the add_partitions() API to authorize only new partitions, improving performance and reducing authorization overhead.

Apache Jira: [HIVE-28371](#)

CDPD-73046: Removal of duplicated proto reader/writer classes

7.3.1.500

Duplicate Java files for proto reader/writer classes were present in Hive, which were already available in Apache Tez. These duplicates caused redundancy and missed improvements introduced in Tez, such as those from TEZ-4296, TEZ-4105, and TEZ-4305.

The issue was fixed by removing the duplicated proto reader/writer classes from Hive, ensuring the use of the improved versions available in Apache Tez.

Apache Jira: [HIVE-28028](#)

CDPD-74539: Maria DB falls back to MySQL in Hive

7.3.1.500

Hive downstream had errors in supporting Maria DB.

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

CDPD-77713: Deadlock occurs in TxnStoreMutex when acquiring lock

7.3.1.500

Deadlocks occurred in Hive Metastore due to MySQL's REPEATABLE-READ isolation level, which caused locking conflicts during housekeeping tasks.

The issue was addressed by restoring the TxnHandler's isolation level to READ-COMMITTED.

Apache Jira: [HIVE-28669](#)

CDPD-77905: MRCompactor causes data loss during major compaction

7.3.1.500

During a major compaction, records matching certain conditions were lost due to incorrect handling in MRCompactor.

The issue was addressed by ensuring that all records are correctly preserved during major compaction.

Apache Jira: [HIVE-28700](#)

CDPD-75656: OOM when compiling query with many GROUP BY columns aliased multiple times

7.3.1.500

HiveServer2 became unresponsive and crashed with an OutOfMemoryError when compiling queries that include GROUP BY columns aliased multiple times in the SELECT clause.

The issue was addressed by customizing the metadata handler to limit the growth of unique key derivation.

Apache Jira: [HIVE-28582](#)

DWX-20754: Error while running LATERAL VIEW query on non-native tables

7.3.1.500

When running LATERAL VIEW queries on non-native Iceberg tables, users encountered the error `org.apache.hadoop.hive.ql.parse.SemanticException: Line 0:-1 Invalid column reference 'BLOCK__OFFSET__INSIDE__FILE'`.

This issue occurred because Iceberg tables were incorrectly classified as native tables, which led to the addition of incorrect virtual columns in the RowResolver. This issue has been resolved, and LATERAL VIEW queries on non-native Iceberg tables should now work as expected.

Apache Jira: [HIVE-28938](#)

Cloudera Runtime 7.3.1.400 SP2:

CDPD-81766: Database Setting Consistency in Spark3 HWC

7.3.1.400

Spark3's Hive Warehouse Connector (HWC) did not consistently apply the database setting when validating if a table existed during append mode writes. This led to inconsistencies where the database setting was not used for validation, even though data was correctly written to the intended database.

This issue was resolved and is now in a patch-ready state. This ensures the database setting is consistently applied during table validation in Spark3 HWC, preventing prior inconsistencies.

CDPD-81122: Enhanced Concurrent Access in HWC Secure Mode

7.3.1.400

Spark applications running multiple concurrent queries in HWC's SECURE_ACCESS mode encountered failures and correctness problems. This happened because the system faced difficulties when generating temporary table names and managing staging directories simultaneously for multiple reads.

This issue was addressed by improving the handling of concurrent operations within HWC's SECURE_ACCESS mode.

CDPD-81453: Efficient Handling of Timed-Out Transactions in Replication

7.3.1.400

Hive replication did not log transactions that timed out as 'ABORTED'. This caused these transactions to remain on the target cluster for an extended period.

This issue was resolved by ensuring that transactions aborted due to timeout are now properly logged. This allows their abort event to be replicated, leading to prompt removal from the target environment.

Apache Jira: [HIVE-27797](#)

CDPD-81420: Table Filtering for Ranger Policies

7.3.1.400

Ownership details for tables were not correctly carried through the system during filtering, which prevented Ranger from applying policies based on who owned the tables.

This issue was resolved by ensuring that ownership information is now consistently included when tables are filtered. This allows Ranger to accurately enforce policies based on table ownership, leading to improved performance when filtering databases and tables.

CDPD-77626: Improving performance of ALTER PARTITION operations using direct SQL

7.3.1.400

Running ALTER PARTITION operations using direct SQL failed for some databases. The failures occurred due to missing data type conversions for CLOB and Boolean fields, causing the system to fall back to slower ORM (Object Relational Mapping) paths.

The issue was addressed by adding proper handling for CLOB and Boolean type conversions. With this fix, ALTER PARTITION operations now run successfully using direct SQL.

Apache Jira: [HIVE-28271](#), [HIVE-27530](#)

Cloudera Runtime 7.3.1.300 SP1 CHF 1

CDPD-64950: Deadlock during Spark shutdown due to duplicate transaction cleanup

7.3.1.300

During Spark application shutdown, transactions were being closed by two separate mechanisms at the same time. This parallel cleanup could result in a deadlock, especially when the heartbeat interval was set to a low value.

The issue was addressed by ensuring that transaction cleanup occurs through a single mechanism during shutdown, avoiding concurrent execution and potential deadlocks.

CDPD-78334: Support custom delimiter in SkippingTextInputFormat

7.3.1.300

Queries like SELECT COUNT(*) returned wrong results when a custom record delimiter was used. The input file was read as a single line because the custom delimiter was ignored.

The issue was addressed by ensuring that the custom record delimiter is considered while reading the file, so that queries work as expected.

Apache Jira: [HIVE-27498](#)

CDPD-79237: Hive Metastore schema upgrade fails due to NULL values

7.3.1.300

Upgrading from CDP Private Cloud Base 7.1.7.2052 to 7.1.9.1010 fails during the Hive Metastore schema upgrade. The upgrade script issues the following command:

```
ALTER TABLE "DBS" ALTER COLUMN "TYPE" SET DEFAULT 'NATIVE', ALTER COLUMN "TYPE" SET NOT NULL;
```

This fails because the DBS.TYPE column contains NULL values. These NULLs are introduced by canary databases created by Cloudera Manager, which insert entries in the HMS database without setting the TYPE.

The issue was addressed by ensuring that canary databases created by Cloudera Manager correctly populate the TYPE column in the DBS table, preventing NULL values and allowing the schema upgrade to proceed.

Cloudera Runtime 7.3.1.200 SP1

CDPD-78342/CDPD-72605: Optimized partition authorization in HiveMetaStore to reduce overhead

7.3.1.200

The `add_partitions()` API in HiveMetastore was authorizing both new and existing partitions, leading to unnecessary processing and increased load on the authorization service.

The issue was addressed by modifying the `add_partitions()` API to authorize only new partitions, improving performance and reducing authorization overhead.

CDPD-77990: Upgraded MySQL Connector/J to 8.2.0 to fix CVE-2023-22102

7.3.1.200

The existing MySQL Connector/J version was vulnerable to CVE-2023-22102.

The issue was addressed by upgrading `mysql-connector-j` to version 8.2.0 in `packaging/src/docker/Dockerfile`.

CDPD-62654/CDPD-77985: Hive Metastore now sends a single AlterPartitionEvent for bulk partition updates

7.3.1.200

HiveMetastore previously sent individual `AlterPartitionEvent` for each altered partition, leading to inefficiencies and pressure on the back db.

The issue was addressed by modifying Hive Metastore to send a single `AlterPartitionEvents` containing a list of partitions for bulk updates, `hive.metastore.alterPartitions.notification.v2.enabled` to turn on this feature.

Apache Jira: [HIVE-27746](#)

CDPD-73669: Secondary pool connection starvation caused by updatePartitionColumnStatisticsInBatch API

7.3.1.200

Hive queries intermittently failed with `Connection is not available`, request timed out errors. The issue occurred because the `updatePartitionColumnStatisticsInBatch` method in `ObjectStore` used connections from the secondary pool, which had a pool size of only two, leading to connection starvation.

The fix ensures that the `updatePartitionColumnStatisticsInBatch` API now requests connections from the primary connection pool, preventing connection starvation in the secondary pool.

Apache Jira: [HIVE-28456](#)

CDPD-61676/CDPD-78341: Drop renamed external table fails due to missing update in PART_COL_STATS

7.3.1.200

When `hive.metastore.try.direct.sql.ddl` is set to false, dropping an external partitioned table after renaming it fails due to a foreign key constraint error in the `PART_COL_STATS` table. The table name in `PART_COL_STATS` is not updated during the rename, causing issues during deletion.

The issue was addressed by ensuring that the `PART_COL_STATS` table is updated during the rename operation, making partition column statistics usable after the rename and allowing the table to be dropped successfully.

Apache Jira: [HIVE-27539](#)

CDPD-79469: Selecting data from a bucketed table with a decimal column throws NPE

7.3.1.200

When `hive.tez.bucket.pruning` is enabled, selecting data from a bucketed table with a decimal column type fails with a `NullPointerException`. The issue occurs due to a mismatch in decimal precision and scale while determining the bucket number, causing an overflow and returning null.

The issue was addressed by ensuring that the correct decimal type information is used from the actual field object inspector instead of the default type info, preventing the overflow and NullPointerException.

Apache Jira: [HIVE-28076](#)

CDPD-74095: Connection timeout while inserting Hive partitions due to secondary connection pool limitation

7.3.1.200

Since HIVE-26419, Hive uses a secondary connection pool (size 2) for schema and value generation. However, this pool also handles nontransactional connections, causing the updatePartitionColumnStatisticsInBatch request to fail with a Connection is not available, request timed out error when the pool reaches its limit during slow insert or update operations.

The issue was addressed by ensuring that time-consuming API requests use the primary connection pool instead of the secondary pool, preventing connection exhaustion.

Apache Jira: [HIVE-28456](#)

CDPD-78331: HPLSQL built-in functions fail in insert statement

7.3.1.200

After the HIVE-27492 fix, some HPLSQL built-in functions like trim and lower stopped working in INSERT statements. This happened because UDFs already present in Hive were removed to avoid duplication, but HPLSQL's local and offline modes still required them.

The issue was addressed by restoring the removed UDFs in HPLSQL and fixing related function issues to ensure compatibility in all execution modes.

Apache Jira: [HIVE-28143](#)

CDPD-78343: Syntax error in HPL/SQL error handling

7.3.1.200

In HPL/SQL, setting hplsql.onerror using the SET command resulted in a syntax error because the grammar file (Hplsql.g4) only allowed identifiers without dots (.).

The issue was addressed by updating the grammar to support qualified identifiers, allowing the SET command to accept dot (.) notation.



Note:

Use EXECUTE 'SET <HS2 Config>' or EXECUTE IMMEDIATE 'SET <HS2 Config>' to configure HS2 settings.

```
Example: EXECUTE 'SET hive.merge.split.update=true';
```

Apache Jira: [HIVE-28253](#)

CDPD-78330: HPL/SQL built-in functions like sysdate not working

7.3.1.200

HPL/SQL built-in functions that are not available in Hive, such as sysdate, were failing with a SemanticException when used in queries. Only functions present in both HPL/SQL and Hive were working.

The issue was addressed by modifying the query parsing logic. Now, HPL/SQL built-in functions are executed directly, and only functions also available in Hive are forwarded to Hive for execution.

Apache Jira: [HIVE-27492](#)

CDPD-78345: Signalling CONDITION HANDLER is not working in HPLSQL

7.3.1.200

The user-defined CONDITION HANDLERS in HPLSQL are not being triggered as expected. Instead of running the handlers, the system only logs the conditions, so the handlers aren't available when needed.

The issue was addressed by ensuring that user-defined condition handlers are properly registered and invoked when a SIGNAL statement raises a corresponding condition.

Apache Jira: [HIVE-28215](#)

CDPD-78333: EXECUTE IMMEDIATE throwing ClassCastException in HPL/SQL

7.3.1.200

When executing a select count(*) query, it returns a long value, but HPLSQL expects a string. This mismatch causes the following error:

```
Caused by: java.lang.ClassCastException: class java.lang.Long cannot be cast to class java.lang.String
    at org.apache.hive.service.cli.operation.hplsql.HplSqlQueryExecutor$OperationRowResult.get
```

The issue was addressed by converting the result to a string when the expected type is a string.

Apache Jira: [HIVE-28215](#)

CDPD-79844: EXECUTE IMMEDIATE displaying error despite successful data load

7.3.1.200

Running EXECUTE IMMEDIATE 'LOAD DATA INPATH "/tmp/test.txt" OVERWRITE INTO TABLE test_table' displayed an error on the console, even though the data was successfully loaded into the table. This occurred because HPL/SQL attempted to check the result set metadata after execution, but LOAD DATA queries do not return a result set, leading to a NullPointerException.

The issue was addressed by ensuring that result set metadata is accessed only when a result set is present.

Apache Jira: [HIVE-28766](#)

CDPD-67033: HWC for Spark 3 compatibility with Spark 3.5

7.3.1.200

The Spark 3.5, based on Cloudera on cloud 7.2.18 libraries, caused a failure in the HWC for Spark 3 build. Canary builds indicate that broke compatibility.

The issue was addressed by updating HWC for Spark 3 to align with Spark 3.5 changes and ensuring compatibility with Cloudera on cloud 7.2.18 dependencies

CDPD-80097: Datahub recreation fails due to Hive Metastore schema validation error

7.3.1.200

Datahub recreation on Azure fails because Hive Metastore schema validation cannot retrieve the schema version due to insufficient permissions on the VERSION table.

This issue is now fixed.

Cloudera Runtime 7.3.1.100 CHF 1

CDPD-74456: Spark3 hwc.setDatabase() writes to the correct database

7.3.1.100

When setting the database using hive.setDatabase("DB") and performing CREATE TABLE or write operations with Hive Warehouse Connector (HWC), the operations were executed in a default database. This issue is now resolved and the operations are executed in the correct database.

The issue is now fixed.

CDPD-74373: Timestamp displays incorrectly in Spark HWC with JDBC_READER mode

7.3.1.100

When using Spark HWC with JDBC_READER mode, timestamps were displayed incorrectly. For example, 0001-01-01 00:00:00.0 was interpreted as 0000-12-30 00:00:00.

This issue is addressed by correcting timestamp handling in JDBC_READER mode to ensure accurate representation of timestamps before the Gregorian calendar was adopted.

CDPD-76932: Incorrect query results due to TableScan merge in shared work optimizer

7.3.1.100

During shared work optimization, TableScan operators were merged even when they had different Dynamic Partition Pruning (DPP) parent operators. This caused the filter from the missing DPP operator to be ignored, leading to incorrect query results.

This issue is resolved by modifying the shared work optimizer to check the parents of TableScan operators and skip merging when DPP edges differ.

Apache Jira: [HIVE-26968](#)

CDPD-78115: Thread safety issue in HiveSequenceFileInputFormat

7.3.1.100

Concurrent queries returned incorrect results when query result caching was disabled due to a thread safety issue in HiveSequenceFileInputFormat.

This issue is now resolved and the files are now set in a thread-safe manner to ensure correct query results.

CDPD-78129: Materialized view rebuild failure due to stale locks

7.3.1.100

If a materialized view rebuild is aborted, the lock entry in the materialization_rebuild_locks table is not removed. This prevents subsequent rebuilds of the same materialized view, causing error

```
Error: Error while compiling statement: FAILED: SemanticException
org.apache.hadoop.hive.ql.parse.SemanticException: Another process
is rebuilding the materialized view view_name (state=42000, code=40000)
```

The fix ensures that the materialized view rebuild lock is removed when a rebuild transaction is aborted. The MaterializationRebuildLockHeartbeater now checks the transaction state before heartbeating, allowing outdated locks to be cleaned properly.

Apache Jira: [HIVE-28416](#)

CDPD-78166: Residual operator tree in shared work optimizer causes dynamic partition pruning errors

7.3.1.100

Shared work optimizer left unused operator trees that sent dynamic partition pruning events to non-existent operators. This caused query failures when processing these events, leading to errors in building the physical operator tree.

The issue was addressed by ensuring that any residual unused operator trees are removed during the operator merge process in shared work optimizer, preventing invalid dynamic partition pruning event processing.

Apache Jira: [HIVE-28484](#)

CDPD-78113: Conversion failure from RexLiteral to ExprNode for empty strings

7.3.1.100

Conversion from RexLiteral to ExprNode failed when the literal was an empty string, causing the cost-based optimizer to fail for queries.

The issue was addressed by ensuring that an empty string literal in a filter produces a valid RexNode, preventing cost-based optimizer failures.

Apache Jira: [HIVE-28431](#)

Cloudera Runtime 7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

The DROP DATA CONNECTOR IF NOT EXISTS command incorrectly throws a NoSuchObject Exception 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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

HiveServer became unstable due to an infinite loop during query parsing with UNION operations, causing an out-of-memory error the 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

7.3.1

The Reader in AcidUtils.isRawFormatFile is not being closed after the check is finished. This causes issues when resources on the DFSCClient 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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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 Impala

Review the list of Impala issues that are resolved in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3**CDPD-87152: System error when renaming tables during metadata invalidations**

7.3.1.500

A system error could occur when you renamed tables at the same time as metadata invalidation commands. This was caused by the system unexpectedly encountering a null table object. The issue is now fixed.

Apache Jira: [IMPALA-13998](#)

CDPD-85651: Query execution time limits

7.3.1.500

Queries sometimes timed out unexpectedly because the set execution time limit (`EXEC_TIME_LIMIT_S`) was inaccurately counting time spent in planning and scheduling.

This issue is resolved by ensuring the query execution time limit (`EXEC_TIME_LIMIT_S`) now starts only when the query begins actual execution.

Apache Jira: [IMPALA-14001](#)

CDPD-80939: Missing equivalence conjunct in aggregation node with inline views

7.3.1.500

In queries that include filters on the result of a **UNION** operation, the planner sometimes removed required conjuncts, which caused incorrect query results.

This has now been fixed.



Note: This issue affected only distinct **UNION** operations, not **UNION ALL**.

Apache Jira: [IMPALA-13873](#)

CDPD-83207/CDPD-84628: Impala cluster responsiveness during table renames

7.3.1.500

Performing **ALTER TABLE RENAME** operations caused Impala to hold a critical internal lock for too long, which blocked other DDL/DMLs.

This issue is resolved by ensuring that the critical internal lock is no longer held during long-running external calls initiated by **ALTER TABLE RENAME** operations. This prevents the entire Impala cluster from being blocked, allowing other queries and catalog operations to proceed without interruption.

Apache Jira: [IMPALA-13631](#)

CDPD-83440: Nested loop join rewrites disjunctive subquery incorrectly

7.3.1.500

Queries with a subquery inside an **OR** condition could return incorrect results. The join was rewritten incorrectly, leading to a wrong comparison.

The issue was addressed by skipping the join rewrite when the subquery is inside an **OR** condition.

Apache Jira: [IMPALA-13991](#)

CDPD-81013: Stable table loading on Ozone with empty files

7.3.1.500

Tables with empty files on Ozone failed to load with a `TableLoadingException`.

The issue was addressed by handling **NULL** values correctly, which prevents `TableLoadingException` when empty files are present.

Apache Jira: [IMPALA-14194](#)

CDPD-81905: Missing tables in Top-N memory usage list on Web UI

7.3.1.500

Some tables did not appear in the Top-N Tables list on the /catalog Web UI after metadata loading in On-demand Metadata mode, such as from a **DESCRIBE** command. The memory metrics were only updated when a **FULL** Thrift object was requested.

The fix includes updating memory usage metrics immediately after metadata is loaded for an **HDFS** table, regardless of the Thrift object type requested.

Apache Jira: [IMPALA-13154](#)

CDPD-83441: Inconsistent Row Count Output in Impala-shell

7.3.1.500

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

This issue was resolved by adding a new option in Impala-shell called `--beeswax_compat_num_rows`. When this option is enabled, Impala-shell now prints Fetched 0 row(s) in along with the elapsed time for all Impala commands.

Apache Jira: [IMPALA-13584](#)

CDPD-74168: Stable file system operations

7.3.1.500

Unreadable paths in Java CLASSPATH could lead to crashing during startup without an informative error message.

This was resolved by ignoring unreadable paths in CLASSPATH.

Apache Jira: [IMPALA-13371](#)

Concurrently running UPDATE and DELETE operations can no longer revive deleted rows

7.3.1.500

This has now been fixed.

Apache Jira: [IMPALA-12787](#)

Metadata query problem from Impala

7.3.1.500

When querying metadata from Impala on a cluster with more than a few nodes, the query might fail with a null pointer exception.

The fix addresses issues noticed while querying metadata from Impala on a cluster with more than a few nodes. The change forces Iceberg metadata scanner fragments to be scheduled on the coordinator. A new flag is introduced in the TPlanFragment thrift struct and if the value is true, then the fragment is always scheduled on the coordinator.

Apache Jira: [IMPALA-12809](#)

: Never smallify existing StringValue objects, only new ones during DeepCopy

7.3.1.500

This has now been fixed.

Apache Jira: [IMPALA-13138](#)

Addressing IllegalStateException in Complex Views post upgrade

7.3.1.500

When executing queries that generate runtime filters where the same column identifier appears repeatedly after upgrading to the 2024.0.18.0-206 runtime, you encountered the following error: `IllegalStateException: null`. This issue has been fixed.

This has now been fixed.

Apache Jira: [IMPALA-13270](#)

Stability Improvement for analytic functions on collections

7.3.1.500

The sorting process incorrectly included unnecessary elements, causing errors during array operations and leading to frequent query failures. This issue has been fixed to ensure only complete data entries are used in the sorting process. This prevents crashes and maintains stable execution of analytic functions on collections.

This has now been fixed.

Apache Jira: [IMPALA-13272](#)

CDPD-73187: Impala Ranger audit plugin fails to create audit logs

7.3.1.500

The fix ensures that the Ranger plug-in in Hive and Impala send audit events to the Solr service after upgrading Data Lake to a version that requires SSL for Ranger's audit events.

This has now been fixed.

Apache Jira: [CDPD-73187](#)

IMPALA-13170: Database deletion during metadata fetch

7.3.1.500

Running `SHOW DATABASES` in Impala while simultaneously dropping a database in Hive caused an `InconsistentMetadataFetchException` due to missing database metadata.

Improved exception handling for operations listing databases and tables. If a database is dropped during metadata fetch, the operation now handles the missing database.

Apache Jira: [IMPALA-13170](#)

CDPD-56130: Event Processor failure due to table lock release error

7.3.1.500

The event processor entered an error state when a table write lock attempt timed out. Since the lock was not held by the current thread, an attempt to release it triggered an `IllegalMonitorStateException`.

The fix added a check to ensure the table holds the write lock before releasing it, preventing event processor failures and unnecessary table invalidations.

Apache Jira: [IMPALA-12141](#)

CDPD-70945: ConcurrentModificationException in ReloadEvent processing

7.3.1.500

The event processor encountered a `ConcurrentModificationException` when handling partition-level `RELOAD` events. This occurred because `ReloadEvent.isOlderEvent()` checked the partition reload status without holding a table read lock, leading to conflicts with concurrent DDL/DML operations modifying the partition list.

The issue was addressed by ensuring that the event processor acquires a table read lock before checking partition reload status, preventing concurrent modifications from causing errors.

Apache Jira: [IMPALA-13126](#)

CDPD-72058: Optimized serialization of position delete records

7.3.1.500

The serialization of position delete records was inefficient, storing file paths redundantly and increasing memory usage. The process involved unnecessary copying, leading to larger buffers that required compression before transmission.

The fix includes optimizing serialization by grouping position delete records with the same file path.

Apache Jira: [IMPALA-13194](#)

CDPD-74683: Incorrect skipping of file metadata reload for ALTER_TABLE events

7.3.1.500

The optimization introduced in [IMPALA-12487](#) skipped file metadata reload for `ALTER_TABLE` events if changes in storage descriptor were trivial. However, some HMS clients modify both table properties and storage descriptor. If table properties had non-trivial changes (e.g., a location change), skipping the file metadata reload led to inconsistencies.

The issue was addressed by refining the checks for skipping file metadata reload in `ALTER_TABLE` events.

Apache Jira: [IMPALA-13403](#)

IMPALA-12607: Optimized event fetching by filtering at the metastore

7.3.1.500

Impala fetched all metastore events and filtered them in its cache based on dbName/tableName. This approach became a bottleneck when handling a large number of events.

Impala now directly fetches events specific to the database or table from the metastore by leveraging the HIVE-27499 change.

Apache Jira: [IMPALA-12607](#)

CDPD-82599: Query rejected due to too many fragment instances on one node

7.3.1.500

Some queries failed with a scheduling error when too many fragment instances were placed on a single executor node.

The fix limits the number of instances per node during scheduling to avoid exceeding the maximum allowed.

Apache Jira: [IMPALA-14006](#)

CDPD-82303: EXEC_TIME_LIMIT_S incorrectly includes planning time

7.3.1.500

The EXEC_TIME_LIMIT_S timer was triggered during the planning and admission (queuing) phases, which could cause queries to fail before any processing on the backend started.

The issue was addressed by starting the EXEC_TIME_LIMIT_S only after the query is ready to run on the backend. This ensures the timer applies only to the actual processing phase.

Apache Jira: [IMPALA-14001](#)

CDPD-82364: Nested loop join rewrites disjunctive subquery incorrectly

7.3.1.500

Queries with a subquery inside an OR condition could return incorrect results due to an improper join predicate rewrite. The issue was addressed by skipping the join rewrite when the subquery is part of a disjunctive (OR) expression.

This has now been fixed.

Apache Jira: [IMPALA-13991](#)

Invalid cardinality calculation in sortnode's computestats

7.3.1.500

An error occurred during query execution due to an overflow in the calculation of row limits, causing unexpected failures. The calculation was updated to prevent overflow and eliminate the error.

This has now been fixed.

Apache Jira: [IMPALA-14070](#)

DWX-18050: Ranger audit logs show origin client's IP address for Impala

7.3.1.500

Ranger audit logs now show the origin client's IP address for Impala when Impala coordinator's config has use_xff_address_as_origin=true. This applies to all Impala clients such as impala-shell, impyla, jdbc and odbc clients.

Workaround for Hue: When using Hue client, enable_xff_for_hive_impala=true in hue-safety-valve configuration is additionally also required to get origin client's IP in Ranger audit logs.



Note: This change does not fix the issue where Cloudera Data Visualization connections log an internal IP address and not the origin client's IP address in Ranger audit.

Apache Jira: [IMPALA-13312](#)

Cloudera Runtime 7.3.1.400 SP2**CDPD-84214: User authentication consistency for Impala queries**

Impala queries previously failed with an authorization error, indicating a mismatch between the user authorized on the Knox connection and the session username. This issue occurred intermittently, typically on the first attempt when running Iceberg queries, especially when cookie-based authentication was used after an initial secure authentication setup.

This issue was resolved by ensuring that user identity is correctly maintained during cookie-based authentication, even after initial secure setup. This prevents mismatches in user verification, allowing queries to proceed without authorization errors.

Apache Jira: [IMPALA-14083](#)

Cloudera Runtime 7.3.1.300 SP1 CHF 1**CDPD-68079: Catalog Object Cache Size Calculation in JDK 11+**

7.3.1.300

An error occurred in JDK 11+ where the catalog object cache incorrectly calculated the size of an object, resulting in an exception message:

```
java.lang.reflect.InaccessibleObjectException: Unable to make field private jdk.internal.platform.cgroupv1.CgroupV1SubsystemController jdk.internal.platform.cgroupv1.CgroupV1Subsystem.cpu accessible.
```

The issue has been resolved.

CDPD-80169: File metadata reload incorrectly skipped on ALTER TABLE

7.3.1.300

Impala skipped reloading file metadata on ALTER TABLE events if changes in the StorageDescriptor were trivial, even when table properties had important changes like a new location.

The fix includes updating the logic to ensure file metadata reload is not skipped when table properties have non-trivial changes.

Apache Jira: [IMPALA-13403](#)

CDPD-80168: Missing tables in Top-N memory usage list on Web UI

7.3.1.300

Some tables did not appear in the Top-N Tables list on the /catalog Web UI after metadata loading, such as from a DESCRIBE command. The memory metrics were only updated when a FULL Thrift object was requested.

The fix includes updating memory usage metrics immediately after metadata is loaded for an HDFS table, regardless of the Thrift object type requested.

Apache Jira: [IMPALA-13154](#)

CDPD-80170: Inconsistent ACID base folder validation between Impala and Hive

7.3.1.300

7.3.1.300

Impala rejected base folders created by INSERT OVERWRITE or TRUNCATE if earlier write IDs were still open, causing read inconsistencies with Hive.

Updated logic to accept these base folders, matching Hive behavior and ensuring consistent reads across both engines.

Apache Jira: [IMPALA-13759](#)

Cloudera Runtime 7.3.1.200 SP1**CDPD-80166: Ignore CREATE_TABLE events for inaccessible databases to prevent event processor error**

7.3.1.200

CREATE_TABLE events for databases restricted by authorization caused the Impala event processor to enter an ERROR state

The event processor now ignores such events when the database is not found, logging the issue for administrators.

Apache Jira: [IMPALA-11735](#)

CDPD-78277: Customize timezone for UNIXTIME_MICROS columns in Kudu

7.3.1.200

Impala used the server's timezone for UNIXTIME_MICROS columns, causing mismatches with Spark, which uses UTC.

Added an option to set the timezone, ensuring consistent timestamps between Impala and Spark.

Apache Jira: [IMPALA-12370](#)

CDPD-79017: Fixed lost exceptions during re-analysis failures

7.3.1.200

When an AnalysisException occurs during re-analysis, calling toSql() throws another exception, causing the original exception to be lost.

The issue was addressed by ensuring the original exception is not lost when an error occurs during query re-analysis. Now, the error message properly logs the root cause instead of being replaced by another exception.

Apache Jira: [IMPALA-12811](#)

Cloudera Runtime 7.3.1.100 CHF 1**CDPD-78366: Session management for Hue logins with Impala**

7.3.1.100

When logging into the Hue web UI, users occasionally encountered the following error:

```
The user authorized on the connection 'hue/gateway0.xyz.site@XYZ.SITE' does not match the session username 'hue/gateway1.xyz.site@XYZ.SITE'
```

This issue arised when a Hue user's request was routed to a different Hue backend other than the one that initiated the original HS2 session with Impala. With Kerberos authentication, Impala restricts session reuse across different hosts or realms, even if the user remains the same.

To address this, proxy clients like Hue can now reuse HS2 sessions across multiple hosts or realms. This is allowed as these proxy clients can be trusted. This improvement enhances session management and ensures a smoother user experience when using Hue.

Apache Jira: [IMPALA-11298](#)

CDPD-78368: Executor crash during runtime filter generation

7.3.1.100

Enabling the MIN_MAX runtime filter in nested loop Joins for certain queries caused out-of-bounds access, leading to executor crashes that prevented any queries from running, disrupting production.

The issue is resolved by ensuring ScalarExprEvaluator properly calls the Open() function, preventing out-of-bounds access and stabilizing query execution.

Apache Jira: [IMPALA-12582](#)

CDPD-78364: Impalad crash due to inconsistent tuple IDs during query execution

7.3.1.100

Avoid impalad crash in RowDescriptor::InitTupleIdxMap() by verifying tuple ids in descriptor table received in executor side.

Apache Jira: [IMPALA-13378](#)

CDPD-78170: Boolean literals in OR conditions are not simplified as expected

7.3.1.100

Impala currently fails to simplify expressions like "id = 0 OR false" to "id = 0". The root cause is that the CompoundPredicate generated by NormalizeExprsRule is not analyzed, preventing SimplifyConditionalsRule from applying the expected rewrite.

This fix addresses the issue by ensuring that the rewritten CompoundPredicate is analyzed properly, allowing boolean literals in AND/OR conditions to be simplified as expected.

Apache Jira: [IMPALA-13203](#)

CDPD-78177: Conjunct registration issue during query rewrite in Impala

7.3.1.100

This issue, introduced in version Cloudera Base on premises 7.1.9, caused some queries with expressions that could be rewritten to false to fail with an IllegalStateException: Illegal reference to non-materialized slot

The issue was addressed by reverting an optimization that skipped registering certain conjuncts, ensuring proper registration and assignment during query analysis.

Apache Jira: [IMPALA-13302](#)

CDPD-78353: Failed table loads not retried after metastore recovery

7.3.1.100

Tables failed to load when the metastore was down, and queries continued to fail even after the metastore was back up, requiring manual invalidation.

The system now automatically retries loading failed tables when a query is executed, ensuring successful loads once the metastore is up, without manual intervention.

Apache Jira: [IMPALA-13120](#)

CDPD-78175: Incorrect results due to predicate pushdown in join condition

7.3.1.100

Impala can incorrectly push predicates to scan nodes, causing wrong results in some join scenarios. For example, a query that should return 0 rows may return incorrect results when specific predicates are used.

Removing analytic predicates with self-referencing TupleIds during query execution prevents incorrect pushdown and resolves the issue.

Apache Jira: [IMPALA-13262](#)

CDPD-78351: Support x5c Parameter in JSON Web Keys (JWK)

7.3.1.100

Impala coordinator failed to parse in JSON Web Key Set (JWK) and failed to start the coordinator. This issue is now resolved.

Apache Jira: [IMPALA-12559](#)

CDPD-78564: Partition value mismatch during INSERT event processing

7.3.1.100

Impala failed to process some INSERT events due to partition value mismatches. Hive encodes partition strings in file paths using URL encoding, but partition strings in HMS events are not encoded. Impala mismatched partitions during HMS events processing by decoding them incorrectly, causing event processing failures.

The fix includes decoding partition strings only when derived from file paths and not from HMS events to ensure accurate partition mapping.

Apache Jira: [IMPALA-13691](#)

CDPD-78559: Prometheus metric name incompatibility with JDK11 and JDK17

7.3.1.100

Impala metrics derived from JVM MemoryPoolMXBean names contain apostrophe characters in JDK11 and JDK17. Since Prometheus does not support apostrophes in metric names, these metrics cannot be consumed by Prometheus.

To address the issue, apostrophes in Prometheus metric names are now automatically translated to underscores, ensuring compatibility with Prometheus requirements.

Apache Jira: [IMPALA-13638](#)

CDPD-78161: Database deletion during metadata fetch

7.3.1.100

Running SHOW DATABASES in Impala while simultaneously dropping a database in Hive caused an InconsistentMetadataFetchException due to missing database metadata.

Improved exception handling for operations listing databases and tables. If a database is dropped during metadata fetch, the operation now handles the missing database.

Apache Jira: [IMPALA-13170](#)

CDPD-78164: Prevent Impala executor crash on restart

7.3.1.100

Impala crashed when executors received queries with zero fragment instances due to network issues.

Executors now ignore queries with zero fragment instances, and coordinators ensure they do not send such queries.

Apache Jira: [IMPALA-13107](#)

CDPD-78363: DDL hangs with SYNC_DDL when Catalogd switches to standby

7.3.1.100

When Catalogd changes from active to standby while waiting for SYNC_DDL version, it stops receiving catalog topic updates from the statestore. This caused DDL queries to hang indefinitely.

To address the issue, Catalogd now regenerates its service ID when changing to standby and throws an exception if the service ID changes while waiting for SYNC_DDL version.

Apache Jira: [IMPALA-13134](#)

CDPD-78365: Queries canceled after statestore failover

7.3.1.100

Queries were canceled after a statestore failover due to temporary inconsistencies. While a grace period was already in place for statestore restarts, failovers were not handled the same way.

The fix applies the existing post-recovery grace period to statestore failovers, preventing query cancellations caused by temporary disruptions.

Apache Jira: [IMPALA-13159](#)

CDPD-78563: Configurable TCP keepalive for client connections

7.3.1.100

Client connections get lost without an explicit close due to machine resets, network disruptions, or load balancer idle timeouts, leading to resource leaks in Impala.

To address this, new startup options allow you to enable and configure TCP keepalive settings for client connections. These options help detect and close dead connections, preventing premature disconnections.

Apache Jira: [IMPALA-13253](#)

CDPD-78162: Handling invalid column stats in table loading

7.3.1.100

Tables with illegal column stats failed to load, preventing operations like DROP STATS or DROP TABLE.

Impala now normalizes invalid stats to -1 and logs warnings, allowing table recovery.

Apache Jira: [IMPALA-13102](#)

Cloudera Runtime 7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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 Hue

Review the list of Hue issues that are resolved in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

CDPD-57725: Security Enhancement: CSP Nonces Enforced in Hue

Hue now enforces Content Security Policy (CSP) nonces, eliminating unsafe-inline JavaScript to improve security and mitigate InfoSec scan issues related to wildcard domains and other directives.

Update the following in the hue.ini configuration:

```
[desktop]
csp_nonce=true
```



Note: The Workflow Editor will not be available in this mode. This update enhances security by reducing the risk of cross-site scripting (XSS) attacks.

CDPD-81753: Added a configurable flag to optionally reenale data preview on database views in Hue

Previously, data preview for database views was disabled by default due to resource strains caused by complex or long-running views, impacting data validation and analysis.

This issue has now been improved by adding the new `allow_sample_data_from_views` flag, with the default value set to false. Setting this flag to true enables Hue to fetch sample data for database views and thus restore the data preview functionality in the SQL assist panel. You can enable the flag by performing the following steps:

1. Navigate to Cloudera Manager Clusters Hue service Configuration .
2. In the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` field, specify the following parameter:

```
[metastore]
allow_sample_data_from_views=true
```

3. Click Save Changes.

Cloudera Runtime 7.3.1.400 SP2:

There are no fixed issues identified in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

CDPD-80268: Hide logout button when configured with Knox

7.3.1.300

In environments where Hue is configured with Knox, the Hue interface displays the logout button, although logout must be handled by Knox. The logout button is now hidden in the Hue interface when Knox authentication is enabled.

CDPD-77746: Metrics Collection Failure with "DPI-1010: not connected" on Oracle DB backend

7.3.1.300

Metrics collection fails with the error "DPI-1010: not connected" when using an Oracle database as the backend.

This issue is now fixed.

CDPD-76594: Impala Iceberg Table Creation Fails in Hue

7.3.1.300

Previously, creating Iceberg tables in Impala using the Hue importer failed. This issue has now been resolved, allowing seamless table creation.

Cloudera Runtime 7.3.1.200 SP1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.100 CHF 1**CDPD-72754: Enhanced AI Integration in Hue SQL AI Assistant**

7.3.1.100

The Hue SQL AI Assistant now supports Cloudera AI Workbench, Cloudera AI Inference service, and vLLM. These integrations enhance the Hue SQL AI Assistant by enabling the use of private models hosted within Cloudera-managed infrastructure. This ensures enhanced security and privacy while leveraging GenAI for the Hue SQL-related tasks.

- Cloudera AI Workbench: This enables you to securely deploy and run your own models within a virtual private cloud. This configuration enhances control and privacy within your environment. For more information, see [Configure SQL AI Assistant using Cloudera AI Workbench](#).
- Cloudera AI Inference service: Helps in a production-grade serving environment for hosting predictive and generative AI models. This service simplifies model deployment and maintenance. For more information, see [Configure SQL AI Assistant using Cloudera AI Inference service](#).
- vLLM: Provides a customizable environment for running your own language model server. It ensures security and privacy control within your infrastructure. For more information, see [Configure SQL AI Assistant using vLLM](#).

CDPD-72496: File extension restrictions for Hue file uploads

7.3.1.100

Earlier, Hue permitted uploading all file types to the configured filesystems, including unsupported extensions, which posed a security risk.

To enhance security, Hue now allows restricting specific file extensions across all configured filesystems. For example, you can allow .csv file uploads while blocking .exe files. By default, no file extensions are restricted during file uploads.

For more information, see [Managing file extensions for Hue uploads](#).

CDPD-72032: Searching with the application ID and DAG ID is not working in Hive Queries

7.3.1.100

This fix addresses the issue by ensuring that searching using the application ID and DAG ID in the Hive Queries tab of the Job Browser using the Hue Query Processor retrieves the results.

CDPD-71994: Including search within the document as we used to have in CDH

7.3.1.100

In Cloudera, there was a unified search at the top, and the results were not displayed in a tabular format. And these resultant workflows could be copied. This issue is now resolved, and the search is now similar to CDH, where the search result is displayed in a tabular format, and by selecting the results, the workflows can be copied.

CDPD-72777: It's not possible to search a Hue document with the intention to copy it

7.3.1.100

The copy action was not provided in the top search or the left assist filter. This issue is now resolved.

CDPD-71074, CDPD-69265: Hue startup failure resolved on the FIPS-enabled clusters

In the FIPS-enabled clusters, Hue startup failure occurred due to the use of the pycpg2-binary package, which includes an OpenSSL version incompatible with FIPS requirements. The issue now can be resolved by removing the pycpg2-binary package and installing pycpg2 along with its necessary dependencies.

CDPD-67341: Refactor and improve IDBroker support in Hue

7.3.1.100

Refactored the IDBroker support, and more preference is now given to Ranger Authorization Service (RAZ) when both are configured in Hue. Improved the IDBroker HA code section to switch

over to a healthy instance correctly and not depend only on the first one for every scenario. This fix also improves Hue page loading performance.

CDPD-66706: HDFS file compression in Hue is failing

7.3.1.100

Hue HDFS file compression feature did not work and failed with an error message in the OOOIE application logs. This issue is now resolved.

CDPD-73657: Fix configuration issues in HiveServer2 Zookeeper discovery to handle failover

7.3.1.100

- Used the correct port number based on the transport mode.
- Updated the port number and Kerberos principal in the new active endpoint.
- Sorted HiveServer2 endpoints by Zookeeper sequence number to ensure the active endpoint is selected.

CDPD-64979: CORS is too permissive for the public APIs

7.3.1.100

This issue is resolved by adding the following three properties to control Cross-Origin Resource Sharing (CORS) under the [desktop] section in the Hue configuration file:

- `cors_enabled`: Used to enable or disable CORS. The default value is True.
- `cors_allow_credentials`: Used to determines whether the server allows cookies in the cross-site HTTP requests. The default value is True.
- `cors_allowed_origins`: A comma-separated list of origins allowed for CORS. For example:

```
[desktop]
cors_allowed_origi
ns=[***ORIGIN-1***],[***ORIGIN-2***],[***ORIGIN-3***]
```

If no values are set, then all origins are allowed.

To change the values of these properties, go to Cloudera Manager Clusters Hue Configuration Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini`, and specify the configuration in the [desktop] section with the required value. For example, to disable CORS, specify the following:

```
[desktop]
cors_enabled=false
```



Note: You must enable CORS for using public APIs and authenticating using SAML.

CDPD-64286: Hue - Metrics cannot restore database connections on failure

7.3.1.100

This fix resolves an issue where Hue could not establish a new database connection following a backend database restart, ensuring that all metrics are computed consistently after such events.

CDPD-72775: Copies of shared documents are not visible to the user

7.3.1.100

Fixed an issue where copies of shared documents were not visible to the users.

CDPD-62397: Improved default editor selection in Hue configuration

7.3.1.100

Previously, when you tried to log in and access Hue, the first available editor was used by default if multiple editors were configured. This issue is now resolved, allowing you to specify a default editor in the Hue configuration file for all users.

CDPD-57994: Hue import not creating External tables

7.3.1.100

Previously, there was a restriction to create only managed tables with Parquet and Optimized Row Columnar (ORC) formats. This issue is now resolved, and external tables can now be created.

CDPD-71063: Hue - Upgrade eventlet to 0.35.2 and dnspython to 2.6.1 due to CVE-2023-29483

7.3.1.100

Addressed the security fix for CVE-2023-29483 in Hue by ensuring the DNS resolver waits for a valid packet within the full expected time window, thereby providing protection against this type of attack.

Cloudera Runtime 7.3.1**CDPD-65034: Receiving Error " TypeError: 'NoneType' object is not callable" in TCLIService.py when custom headers are being set**

7.3.1

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 CHF 3 or 7.3.1 versions.

Fixed Issues in Apache Iceberg

Review the list of Iceberg issues that are resolved in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Fixed issues in Cloudera Runtime 7.3.1.500 SP3**CDPD-87901: Hive-Iceberg queries fail with IllegalStateException**

7.3.1.500

When running certain Hive-Iceberg queries, the query could fail with an IllegalStateException and you may notice the following error message:

```
ERROR : FAILED: IllegalStateException Invalid file: file length 0
is less than minimal length of the footer tail
```

This error was caused by a race condition related to table statistics files. A query process would sometimes attempt to read a statistics file immediately after its creation but before any data had been written to it. This resulted in the query trying to read an empty (0-length) file, which is invalid, triggering the exception.

This issue is now resolved.

Apache Jira: [HIVE-28278](#)

CDPD-85718: Hive-Iceberg queries fail with NullPointerException while updating CBO statistics

7.3.1.500

When running certain Hive queries (e.g., TPCB benchmarks) on Iceberg tables, the queries would fail with a java.lang.NullPointerException. This error occurred specifically when the CBO engine attempted to update column statistics for query planning. The root cause was due to a case sensitivity mismatch in column names.

This issue is now resolved with a fix that makes the column lookup process case-insensitive.

Apache Jira: [HIVE-28886](#)

CDPD-85557: Invalid column reference in lateral view queries

7.3.1.500

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

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

This issue is now resolved.

Apache Jira: [HIVE-28938](#)

CDPD-80575: Rebuilding a Hive materialized view stored in Iceberg fails with a NullPointerException

7.3.1.500

If you are rebuilding a Hive materialized view stored in Iceberg, the rebuild operation fails with a `NullPointerException` when the schema contains a `varchar` column.

This issue has now been resolved, and the exception is no longer thrown when rebuilding materialized views with a `varchar`-typed column.

Apache Jira: [HIVE-28817](#)

CDPD-80494: Upgrade Netty to 4.1.119.Final

7.3.1.500

This fix addresses upgrading the Netty library in Iceberg to version 4.1.119.Final to fix multiple security vulnerabilities (CVE-2025-24970, CVE-2025-25193, CVE-2024-47535) and resolves compatibility issues with Apache Arrow.

CDPD-77079: month() partition transform fails in Spark SQL for Iceberg tables

7.3.1.500

When you attempt to add a new partition field to an Iceberg table using the singular `month()` transform, the operation fails and results in an `UnsupportedOperationException`. However, the plural form, `months()` works correctly.

This issue has been resolved. The Spark-Iceberg connector has been updated to correctly recognize and support the `month()` partition transform.

Fixed issues in Cloudera Runtime 7.3.1.400 SP2

CDPD-85165: Slow HMS metadata summary collection for non-native Iceberg tables

7.3.1.400

When using the Hive MetaTool to generate a metadata summary (`MetaToolTaskMetadataSummary`), the process was exceptionally slow for databases containing a large number of non-native tables, particularly Iceberg tables.

This issue has now been addressed by introducing a new configuration property, `hive.metatool.summary.nonnative.threads`. This property allows you to specify the number of threads dedicated to collecting summaries for non-native tables like Iceberg. The default value is set to 20.

Apache Jira: [HIVE-28990](#)

Fixed issues in Cloudera Runtime 7.3.1.300 SP1 CHF1

CDPD-75411: SELECT COUNT query on an Iceberg table in AWS times out

7.3.1.300

In an AWS environment, a `SELECT COUNT` query that is run on an Iceberg table times out because some 4KB ORC file parts cannot be downloaded. This issue occurs because Iceberg uses the positional delete index only if the count of positional deletes are less than a threshold value which is by default, 100000.

This issue has been resolved, and the positional delete index is now always used regardless of the positional delete count, resulting in improved performance.

CDPD-79741: Balance scheduling for consecutive partitions for Iceberg tables

7.3.1.300

During remote read scheduling Impala does the following:

Non-Iceberg tables:

- The scheduler processes the scan ranges in partition key order
- The scheduler selects N executors as candidates
- The scheduler chooses the executor from the candidates based on minimum number of assigned bytes
- Therefore, consecutive partitions are more likely to be assigned to different executors

Iceberg tables:

- The scheduler processes the scan ranges in random order
- The scheduler selects N executors as candidates
- The scheduler chooses the executor from the candidates based on minimum number of assigned bytes
- Therefore, consecutive partitions (by partition key order) are assigned randomly and there is a higher chance of clustering

With this fix, IcebergScanNode orders its file descriptors based on their paths to facilitate a more balanced scheduling for consecutive partitions. This is especially important for queries that prune partitions through runtime filters (due to a JOIN), because it does not matter that we schedule the scan ranges evenly, the scan ranges that survive the runtime filters can still be clustered on certain executors.

Apache JIRA: [IMPALA-12765](#)

CDPD-81311: Unable to query Iceberg tables from Impala

7.3.1.300

After upgrading to Cloudera Runtime 7.3.1.200 or lower versions, you may notice issues while querying Iceberg tables from Impala. An error is reported indicating that the migrated file has unexpected schema or partitioning.

In migrated Iceberg tables, there can be data files with missing field IDs. It is assumed that their schema corresponds to the table schema at the point when the table migration happened, which means field IDs can be generated during runtime. The logic becomes complicated when there are complex types in the table and the table is partitioned. In such cases, some adjustments are required during field ID generation and we verify that the file schema corresponds to the table schema (during migration).

This fix ensures that these adjustments are not needed when the table does not have complex types and therefore schema verification is skipped. As a result, Impala can still read the table if there were some trivial schema changes before migration.

Apache JIRA: [IMPALA-13853](#)

Fixed issues in Cloudera Runtime 7.3.1.200 SP1**CDPD-71365: Support Iceberg 1.3 on Spark 3.5**

7.3.1.200

Cloudera Runtime 7.3.1.200 SP1 introduces support for Apache Spark 3.5.4. The Iceberg support for Spark 3.5 is only available in the upstream Iceberg 1.4, however, Cloudera Runtime 7.3.1.200 SP1 offers Iceberg 1.3.

This was addressed and Cloudera ensures that Iceberg 1.3 is compatible with Spark 3.5.4.

CDPD-81709: Update parquet-avro to 1.15.1 due to CVE-2025-30065

7.3.1.200

Due to CVE-2025-30065, schema parsing in the parquet-avro module of Apache Parquet 1.15.0 and earlier versions allows bad actors to execute arbitrary code.

To avoid this CVE, the parquet-avro module is upgraded to version 1.15.1.

Fixed issues in Cloudera Runtime 7.3.1.100 CHF1**CDPD-75667: Querying an Iceberg table with a TIMESTAMP_LTZ column can result in data loss**

7.3.1.100

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

When Impala changes the TIMESTAMP_LTZ column to TIMESTAMP, it does it by calling `alter_table()` on Hive Metastore (HMS) directly. It provides a Metastore Table object to HMS as the desired state of the table. HMS then persists this table object.

This issue is fixed by avoiding the `alter_table()` call to HMS towards the end of loading the Iceberg table. This avoids the necessity of persisting the schema adjustments that Impala had to do while loading the table.

Apache JIRA: [IMPALA-13484](#)

CDPD-78355: Impala should ignore character case of Iceberg schema elements

7.3.1.100

Impala cannot read Iceberg tables written by Apache Spark that contain schema elements in uppercase or lowercase letters.

Schema is case insensitive in Impala, however, Spark allows creation of schema elements with uppercase or lowercase letters and stores them in the metadata JSON files of Iceberg.

With this fix, Impala invokes `Scan.caseSensitive(boolean caseSensitive)` on the `TableScan` object to set case insensitivity.

Apache JIRA: [IMPALA-13463](#)

CDPD-78362: Schema resolution does not work for migrated partitioned Iceberg tables that have complex types

7.3.1.100

Schema resolution does not work correctly for migrated partitioned Iceberg tables that have complex data types. This fix addresses the field ID generation by taking the number of partitions into account. If none of the partition columns are included in the data file (common scenario), file-level field IDs are adjusted accordingly. You could also come across a scenario where all the partition columns are included in the data files.

However, if some partition columns are included in the data file while other partition columns are not, an error is generated.

Apache JIRA: [IMPALA-13364](#)

CDPD-78540: DELETE statement throws DateTimeParseException when deleting from DAY-partitioned Iceberg tables

7.3.1.100

Due to an issue in `IcebergDeleteSink`, Impala cannot successfully run a DELETE operation on Iceberg tables that are partitioned by time-based transforms (YEAR, MONTH, DAY, HOUR).

This fix addresses the error by adding functions that transforms the partition values to their human-readable representations. This is done in the `IcebergDeleteSink` so that the Catalog-side logic is not affected.

Apache JIRA: [IMPALA-12557](#)

CDPD-78562: Iceberg tables have a large memory footprint in catalog cache

7.3.1.100

This fix clears the GroupContentFiles after they are used.

GroupContentFiles stores the file descriptors in Iceberg's format and is used for creating file descriptors in Impala's format. Once the creation is complete, we do not have to retain the Iceberg ContentFiles. Dropping this can significantly reduce the memory footprint of an Iceberg table.

For example, the memory size of a test Iceberg table containing 110k files was reduced from 140MB to 80MB after cleaning the GroupContentFiles.

Apache JIRA: [IMPALA-11265](#)

Fixed issues in Cloudera Runtime 7.3.1

CDPD-49011: Support incremental materialized view rebuild

7.3.1

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

7.3.1

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, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.400 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.200 SP1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.100 CHF 1

There are no fixed issues for Kafka in this release.

Cloudera Runtime 7.3.1

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

7.3.1

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

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

7.3.1

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

OPSAPS-71258: Kafka, Streams Replication Manager, and Streams Messaging Manager cannot process messages compressed with Zstd or Snappy if /tmp is mounted as noexec

7.3.1

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

7.3.1

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

7.3.1

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, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:**CDPD-85458: Prometheus Metrics Enhancement for Kudu**

Kudu now includes Prometheus metrics by adding tablet-level metrics alongside the existing server-level metrics. This improvement provides you with deeper insights to more effectively diagnose performance issues in busy clusters containing up to 2,000 tablets.

Apache Jira:[KUDU-3563](#).

Cloudera Runtime 7.3.1.400 SP2:**CDPD-82273: Backport KUDU-3661 Ranger policy not honored in Kudu**

Fixed an issue in the Ranger authorization provider that could cause some table privileges to be missing in certain environments. This happened when processing the SELECT privilege, which caused the system to stop checking for additional permissions. The issue was primarily seen on RHEL/CentOS 8 systems due to platform-specific behavior in the underlying system libraries.

CDPD-82275: Run a range-aware cluster rebalance with multiple tables

Previously, when rebalancing a cluster with the Kudu command-line tool, the `--enable_range_rebalancing` flag required the `--tables` flag to specify exactly one table. This fix removes that restriction. You can now pass multiple tables to the `--tables` flag when range rebalancing is enabled. Range-partitioned tables among those specified will be rebalanced with ranges considered, while other tables rebalance as usual. If you do not set the `--tables` flag, all tables in the cluster will be rebalanced.

CDPD-81231: Backport KUDU-3638, disable KUDU-3486 behavior by default

This fix addresses an issue with tombstoned tablets caused by functionality introduced in KUDU-3486. The fix disables this functionality. To re-enable the behavior, adjust the `--tserver_send_tombstoned_tablets_report_interval_secs` flag as needed.

CDPD-81229: Backport KUDU-3647 more robust zlib wrapper code

This fix improves robust error handling for the utility wrappers `zlib::CompressLevel`, `Uncompress`()(). It also resolves a data corruption issue in `Base64Decode()`, where trailing bytes encoded as 'A' were previously discarded.

CDPD-81234: Backport KUDU CVE fixes

This fix addresses recently identified security vulnerabilities. The following third-party components have been upgraded:

- `protobuf`: Upgraded to address CVE-2022-3171.
- `zlib`: Upgraded to address CVE-2022-37434.
- `curl`: Upgraded to address CVE-2024-11053.

Resolved a long-standing bug by fixing a race condition within the `OpDriver`. This fix improves stability and ensures the expected behavior of the driver.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.200 SP1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.100 CHF 1

CDPD-78316: Disable KUDU-3367 behavior by default

This fix addresses an issue with major delta compaction caused by functionality introduced in KUDU-3367. The fix disables the functionality by adjusting the default value of the `-all_delete_op_delta_file_cnt_for_compaction` flag. While the original functionality was intended to optimize Kudu deployments of early versions, the functionality introduced by KUDU-3367 isn't relevant for data generated by Kudu servers in Cloudera Base on premises 7.1.1 and newer versions.

Apache Jira: [KUDU-3619](#).

Cloudera Runtime 7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

Kudu now supports configuring a dedicated Spnego keytab.

Apache Jira: [KUDU-3496](#)

KUDU-3497: Optimize leader rebalancer algorithm

7.3.1

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

7.3.1

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 Cloudera Base on premises 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

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

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

CDPD-81958: Improve cookie security in Knox-proxied web UIs with Secure and HttpOnly attributes

The pac4jCsrfToken cookie has now both Secure and HttpOnly flags in Knox proxied applications, improving the security provided by Knox.

Apache JIRA: [KNOX-3134](#)

Cloudera Runtime 7.3.1.400 SP2

CDPD-8148: Knox UI session timeout is not working with SAML authentication

This issue is resolved by the pac4j.cookie.max.age parameter introduced for the pac4j provider, which Knox uses for SAML authentication. This parameter enforces cookie timeout for the cookies created by the pac4j provider.

To set the pac4j.cookie.max.age parameter, go to Cloudera Manager Knox Configuration , and add the following value to the Knox Simplified Topology Management - SSO Authentication Provider field: federation.param.pac4j.cookie.max.age={value}

Apache JIRA: [KNOX-3077](#)

Cloudera Runtime 7.3.1.300 SP1 CHF 1

CDPD-27801: Knox is missing HSTS header for HTTP 404 responses

7.3.1.300, 7.1.9 CHF8

Resolved an issue where Knox was missing the HTTP Strict-Transport-Security response header (HSTS) in HTTP 404 responses. The global HSTS header can now be configured to be included in all HTTP responses.

To configure the HSTS header, go to Cloudera Manager Knox Configuration , search for the Knox Service Advanced Configuration Snippet (Safety Valve) for conf/gateway-site.xml property, and set the following parameters to true:

- gateway.strict.transport.enabled

- `gateway.strict.transport.option`: (Optional) Use this parameter to specify a timeout value for the HTTPS header. This parameter is applicable only if `gateway.strict.transport.enabled` is set to true.

Apache JIRA: [KNOX-3111](#)

CDPD-73368: Knox token management is not working if Cookie Management is enabled

7.3.1.300

Users can now access the Token Management page from the Knox Gateway UI by using KnoxSSO even if Cookie Management is enabled.

Apache JIRA: [KNOX-3060](#)

CDPD-74843: Logs missing in third-party libraries

7.3.1.300

Resolved an issue where third-party libraries had missing logs due to a missing `log4j` library, which affected the ability to diagnose and troubleshoot issues.

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

7.1.9 CHF7, 7.3.1.300

Resolved an issue where the health test for Knox Gateway failed if the `gateway.client.auth.needed` parameter was set to true.

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

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

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

Cloudera Runtime 7.3.1.200 SP1

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

7.3.1.200

Fixed an issue where the Knox Token API raised an `UnknownTokenException` error if the lifespan value of Knox Token TTL was set to -1.

Apache JIRA: [KNOX-3075](#)

CDPD-79963: Knox service might fail due to JARs picked up from the `/usr/share/java` folder

7.3.1.200

Knox service might fail due to Java Archive (JAR) files picked up from the `/usr/share/java` folder.

This issue is now fixed.

Apache JIRA: [KNOX-3108](#)

CDPD-76104: Unable to update the log level for Knox from Cloudera Manager

7.3.1.200

Users were not able to change the log level for Knox from Cloudera Manager. This impacted debugging in case of errors.

This issue is now fixed.

Cloudera Runtime 7.3.1.100 CHF 1

CDPD-74114: Proxyuser groups are not included in POST and PATCH requests

7.3.1.100

Fixed an issue where group headers were not added to POST and PUT requests.

Apache Jira: [KNOX-3062](#)

Cloudera Runtime 7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

You cannot delete custom topologies that were created.

This issue is fixed.

OPSAPS-67480: Default Ranger policy in Cloudera Base on premises 7.1.9 includes cdp-proxy-token topology for new installations, but upgrades do not add cdp-proxy-token to cm_knox policies automatically.

7.3.1

This issue is fixed now.

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

7.3.1

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

This section lists the issues in Livy that are fixed in Cloudera Runtime 7.3.1 release, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.400 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.200 SP1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.100 CHF 1

There are no new fixed issues in this release.

Cloudera Runtime 7.3.1**OPSAPS-71873 - UCL | CKP4| livyfoo0 kms proxy user is not allowed to access HDFS in 7.3.1.0**

Cloudera Manager 7.13.1.0

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

CDPD-73324 - LIVY_FOR_SPARK3 goes into down with Invalid Keystore format error in FIPS cluster 7.3.1

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.

Fixed Issues in Navigator Encrypt

Review the list of Navigator Encrypt issues that are resolved in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.400 SP2**CDPD-82688: Error when building the kernel module**

Navigator Encrypt failed to install on RHEL 7.9 with the following error: Failed installation of navencryptfs 7.1.9.1005 DKMS kernel module!

This issue is fixed now.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.200 SP1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.100 CHF 1

CDPD-77529: Navencrypt 7.1.9 SP1 install failed on Ubuntu 20 - Kernel 5.4.0-200

Depending on the version of Linux kernel module that was running, the NavEncrypt kernel module would fail to install because it would build incorrectly. The error message looked as follows:

```
kernel-deps.h: In function 'zncrypt_get_task_cred': kernel-  
deps.h:45:32: error: passing argument 1 of 'atomic_inc_not_zero'  
from incompatible pointer type [-Werror=incompatible-pointer-  
types] 45 | } while (!atomic_inc_not_zero(&((struct cred *)cred)-  
>usage));
```

This issue is now resolved.

Cloudera Runtime 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 e  
xist 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, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.400 SP2

CDPD-84224: test_oracle_import_with_arithmetic_exception fails due to missing error message in YARN log

7.3.1.400

The Sqoop binary was not correctly updated on the Oozie sharelib, hence the pre-upgrade version of the Sqoop was used, which did not contain the necessary log message. After resolving that issue, the affected executions are now successful.

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

7.3.1.400

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

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

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.200 SP1**CDPD-78069: Oozie action configuration's Java opts are not applied after CDPD-60551**

7.3.1.200

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

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

Cloudera Runtime 7.3.1.100 CHF 1**CDPD-68425: Oozie should validate the backend database identifiers in a case-insensitive way**

7.3.1.100

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

CDPD-76135: Schema check database connection throws PSQLErrorException

7.3.1.100

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

Cloudera Runtime 7.3.1**CDPD-70422: Cannot enforce Oozie parameter `oozie.http.hostname`**

7.3.1

A new property named `oozie.http.bind.hostname` is now introduced to specify the interface that the Oozie Server must be using. This overrides `oozie.http.hostname` set by Cloudera Manager for both HTTP and HTTPS.

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

7.3.1

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

7.3.1

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

7.3.1

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, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3**CDPD-80742: ConstraintViolationException was crashing the ContainerHealthTask in Ozone Recon**

7.3.1.500

The ConstraintViolationException was crashing the ContainerHealthTask task in Recon. This issue is now fixed and the task will not crash and continue to identify the unhealthy containers in SCM.

Apache Jira: [HDDS-12585](#)

CDPD-78266: Ozone Manager displays NullPointerException (NPE) when overwriting empty file using multipart upload

7.3.1.500

Previously, the S3MultipartUploadCompleteRequestWithFSO class displayed the NullPointerException error when an empty file was being overwritten by non-zero file because the omBucketInfo allows the null value. The null value is passed when no update was needed. This issue is now fixed by checking the omBucketInfo before use.

Apache Jira: [HDDS-12131](#)

CDPD-80791: Ozone Recon - Unknown tar ball cleanup for Recon Ozone Manager DB snapshot

7.3.1.500

Previously, when Recon started or restarted, Recon checked for an existing Ozone Manager DB snapshot file if an old tar file was left in the Recon Ozone Manager DB directory location. The old tar file could be left forever, occupying disk space and requiring cleanup if the tarball is large in size. This issue is now fixed by cleaning up any unused or leftover tar files.

Apache Jira: [HDDS-12601](#)

CDPD-81786: Auto-compact tables which can tend to be large in size at intervals

7.3.1.500

Previously, when too many files and directories are deleted in a short period, RocksDB may accumulate too many tombstone which impedes the performance of regular file system operations. This issue is now fixed. The fix adds a background thread at Ozone Manager that can be optionally enabled to compact RocksDB tombstones at a fixed interval by adding the following configuration properties:

Configuration property	Default value	Description
ozone.compaction.service.enabled	false	Enables or disables a background job that periodically compacts rocksdb tables flagged for compaction.
ozone.om.compaction.service.run.interval	6h	Specifies the interval for a background job to periodically compact rocksdb tables flagged for compaction. Units can be defined with postfix and can be ns, ms, s, m, h, or d.
ozone.om.compaction.service.timeout	10m	Specifies a timeout value for the compaction service. If this is set to greater than 0, the service stops waiting for compaction completion after the set time. Units can be defined with postfix and can be ns, ms, s, m, h, or d.
ozone.om.compaction.service.columnfamilies	keyTable,fileTable,directoryTable,deletedTable,deletedDirectoryTable,multipartInfoTable	Specifies a comma separated, no spaces list of all the column families that are compacted by the compaction service. If this is empty, no column families are compacted.

Apache Jira: [HDDS-12819](#)

CDPD-87719: Ozone snapshot diff command is failing

7.3.1.500

The snapdiff operation could fail with a NullPointerException error if a directory had metadata changes, such as Access Control List (ACL) updates or modification time updates between snapshots. This issue is now resolved.

CDPD-83223: Validate Container Balancer move timeout and replication timeout configurations

7.3.1.500

Previously, if the `move.timeout - move.replication.timeout - event.timeout.datanode.offset` configuration was less than 9 minutes, the expiry time set on the delete commands was set in the past. Hence, those delete commands would be invalid. This issue is now fixed by adding the `move.timeout - move.replication.timeout - event.timeout.datanode.offset` rule for container balancer configurations. The timeout value must at least be 9 minutes.

Apache Jira: [HDDS-13068](#)

CDPD-83224: Container Balancer delete commands are sent with an expiration time in the past

7.3.1.500

Previously, delete commands were getting discarded in the container balancer due to incorrect expiration time. This issue is now fixed and the container balancer can now send delete commands correctly.

Apache Jira: [HDDS-13067](#)

Improvements for large scale deletion

7.3.1.500

In this release, to speed up Ozone data deletion, multiple issues are handled. The fixes under this issue focus on several critical areas, including the following areas:

- Performance Improvements
Improving speed that space is reclaimed from the system. Optimization of deletion processes to handle large-scale operations more efficiently.
- Bug Fixes
Addressing issues that occur specifically when deleting large amounts of data.
- Observability Enhancements
Improving observability throughout the deletion process through logs, metrics, and dashboards. New metrics have been introduced and added to Grafana, with a lightweight dashboard created to track the deletion progress.

Apache Jira: [HDDS-11633](#), [HDDS-13057](#), [HDDS-13103](#), [HDDS-13142](#), [HDDS-13259](#), [HDDS-12474](#), [HDDS-11512](#), [HDDS-11511](#), [HDDS-11509](#), [HDDS-11711](#), [HDDS-11779](#), [HDDS-12442](#)

Cloudera Runtime 7.3.1.400 SP2

CDPD-82201: OMKeyAclRequestWithFSO is incorrectly setting full path as key name

7.3.1.400

When you set, add, or remove an ACL for a FSO bucket, the key name gets corrupted with the full key path. This fix ensures the correct key name is set during the ACL calls.

Apache Jira: [HDDS-12891](#)

CDPD-81939: Volume scanner should fail volume if rocksDB is inaccessible

7.3.1.400

When RocksDB becomes unreadable on a DataNode due to disk-related issues, the DataNode will mark the affected storage volume as unhealthy. This proactive health marking enables the system to initiate data replication processes more rapidly, thereby maintaining data availability and integrity.

Apache Jira: [HDDS-12723](#)

CDPD-78932: Container replication should be atomic

7.3.1.400

During container replication, the destination node imports the container from the source node. If any issues are encountered during the import process, the Datanode is responsible for gracefully cleaning up any residual or stale container metadata to maintain system integrity.

Apache Jira: [HDDS-12233](#)

CDPD-73278: Update OM, SCM, Datanode conf for RATIS-2135

7.3.1.400

Set `raft.grpc.message.size.max` to be 1MB larger than `raft.server.log.appender.buffer.byte-limit` for OM, SCM and Datanode.

Apache Jira: [HDDS-11320](#)

CDPD-57559: New Ozone Manager leader cannot verify the Ozone delegation token signed by old Ozone Manager leader

7.3.1.400

If an Ozone cluster is upgraded and then later downgraded, then the new Ozone Manager after downgrade cannot verify the new Ozone delegation tokens issued before downgrade causing the clients which are still running during this upgrade download period to fail. If there is no downgrade, then everything is fine.

This issue is fixed and this fix changes the Ozone delegation token sign from using asymmetric key to symmetric key.

Apache Jira: [HDDS-8829](#)

CDPD-70409: Recon Overview Page UI fails to load if Recon Solr Health throws error

7.3.1.400

This fixes an issue where the Recon UI failed to load if the Solr Health check API threw an error.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

CDPD-80823: Snapshot creation is removing extra keys from the Active Object Storage's DB

7.3.1.300

Wrong keys were trapped in the DeletedTable of the snapshot if the OBS bucket name is a prefix for another OBS bucket, resulting in orphaned blocks. After the fix, it will not remove any extra keys from the DeletedTable for Active Object storage.

Apache Jira: [HDDS-12611](#)

CDPD-73375: Publishing hadoop metrics immediately in Prometheus sink fills up SinkQueue quickly

7.3.1.300

Prometheus sink already have a mechanism to publish metrics every 10 seconds by default using call back with timer event. So, we removed the code to publish immediately and this issue is fixed.

Apache Jira: [HDDS-12193](#)

CDPD-78671: Metric timer task is blocking installSnapshotFromLeader on follower node

7.3.1.300

`ozone.om.snapshot.rocksdb.metrics.enabled` is now available in Ozone to disable metric collection on the snapshotted DB if necessary.

Apache Jira: [HDDS-11339](#)

CDPD-78781: Tarball creation interfering with snapshot purge

7.3.1.300

Synchronize SnapshotDeletingService on BootstrapStateHandler.Lock to make sure that no background service is running when Tarball creation is in progress.

Apache Jira: [HDDS-12210](#)

Cloudera Runtime 7.3.1.200 SP1

CDPD-74556: EC Checksum throws IllegalArgumentException because the buffer limit is negative

7.3.1.200

When ozone.client.bytes.per.checksum is set to a lower value (for example, 16kb), the parity checksum calculation during validation phase is wrong leading to IllegalArgumentException on client. This is now fixed.

Apache Jira: [HDDS-11482](#)

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

7.3.1.200

Default native ACL created for an object such as volume, bucket, or file limits to the object owner and owner's primary group.

Apache Jira: [HDDS-11656](#)

CDPD-72782: Ozone write does not work when http proxy is set for the JVM

7.3.1.200

GRPC uses HTTP internally for its connections and due to this, if HTTP proxy is configured for any Ozone process using GRPC, it directs each call through the proxy even for GRPC which is not desirable for performance. Hence this fix disables proxy for GRPC connections that Ozone uses.

Apache Jira: [HDDS-11257](#)

CDPD-65714: Allow FS client to specify EC as default filesystem replication

7.3.1.200

This fix allows you to specify EC as the default replication type for a file uploaded through the Hadoop FileSystem API to Ozone through a client side configuration option.

Apache Jira: [HDDS-10336](#)

Cloudera Runtime 7.3.1.100 CHF 1

There are no fixed issues in this release.

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.

7.3.1

This issue is now resolved.

OPSAPS-70288: Improvements in master node decommissioning.

7.3.1

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

7.3.1

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.

7.3.1

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.

7.3.1

This issue is now resolved.

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

7.3.1

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.

7.3.1

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.

7.3.1

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.

7.3.1

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.

7.3.1

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.

7.3.1

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.

7.3.1

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.

7.3.1

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.

7.3.1

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.

7.3.1

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.

7.3.1

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 Cloudera 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.

7.3.1

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.

7.3.1

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.

7.3.1

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.

7.3.1

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.

7.3.1

`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.

7.3.1

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.

7.3.1

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.

7.3.1

This issue is now resolved.

Apache Jira: [HDDS-9556](#)

CDPD-62931: Incorrect pipeline ID for closed container.

7.3.1

This issue is now resolved.

Apache Jira: [HDDS-9544](#)

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

7.3.1

This issue is now resolved.

Apache Jira: [HDDS-9542](#)

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

7.3.1

This issue is now resolved.

Apache Jira: [HDDS-9465](#)

CDPD-62466: Improve thread names in Recon.

7.3.1

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.

7.3.1

This issue is now resolved.

Apache Jira: [HDDS-9356](#)

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

7.3.1

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.

7.3.1

This issue is now resolved.

Apache Jira: [HDDS-4368](#)

Fixed Issues in Apache Parquet

Review the list of Parquet issues that are resolved in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.400 SP2

There are no fixed issues in this release.

Fixed issues in Cloudera Runtime 7.3.1.300 SP1 CHF1

There are no new fixed issues in this release.

Fixed issues in Cloudera Runtime 7.3.1.200 SP1

CDPD-81756: Restrict trusted packages in the parquet-avro module

7.3.1.200

Due to CVE-2025-30065, schema parsing in the parquet-avro module of Apache Parquet 1.15.0 and earlier versions allows bad actors to execute arbitrary code.

To prevent this CVE, users must specify all the trusted packages in the `org.apache.parquet.avro.SERIALIZABLE_PACKAGES` environment variable. If the user does not want to specify the override property, then the following packages that are trusted by default are allowed — `java.lang`, `java.math`, `java.io`, `java.net`, `org.apache.parquet.avro`.

Fixed issues in Cloudera Runtime 7.3.1.100 CHF1

There are no fixed issues in this release.

Fixed issues in Cloudera Runtime 7.3.1

There are no fixed issues in this release.

Fixed Issues in Phoenix

Review the list of Phoenix issues that are resolved in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

CDPD-83528: Phoenix salted table range scan issue

7.3.1.500

During a range scan operation on a salted table, data residing in the final salt bucket may not be fully scanned. This issue occurs under specific conditions related to region splits, affecting the salted tables configured for distributed data storage to optimize performance.

This issue is resolved now.

Apache Jira: [PHOENIX-7580](#)

Cloudera Runtime 7.3.1.400 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

CDPD-77685: In Apache Phoenix, a mixed-cased alias does not work in the select statement of INNER JOIN

7.3.1.300

This issue is resolved now.

Apache Jira: [PHOENIX-7491](#)

CDPD-77716: In Apache Phoenix, subqueries that involve both JOIN and UNION fail to work correctly when one side of the UNION yields no results

7.3.1.300

This issue is now resolved by clearing the dependency map, once all the values within it are marked as closed.

Apache Jira: [PHOENIX-7492](#)

Cloudera Runtime 7.3.1.200 SP1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.100 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1

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

7.3.1

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

7.3.1

The OMID update issue is fixed.

OPSAPS-69838: OMID cannot connect contact the ZooKeeper cluster

7.3.1

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

7.3.1

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

7.3.1

This issue is fixed.

Apache Jira: [PHOENIX-6687](#)

CDPD-66730: Phoenix third party upgrade

7.3.1

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 Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

CDPD-83255: Ranger docker support for Knox, RAZ, and Atlas

Ranger docker support has been added for the following services: Knox, RAZ, and Atlas.

Cloudera Runtime 7.3.1.400 SP2

CDPD-78680: Selected long string values getting truncated in the react-select control

In Ranger React UI, the react-select input values are getting truncated for long values.

This issue is fixed. The selected long input string values are not getting truncated on Ranger React UI.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

CDPD-80492: Tomcat upgrade to 9.0.99+

Upgraded Tomcat to 9.0.99+ to address CVE-2025-24813.

CDPD-80018: Concurrent policy label update leads to an infinite loop

Concurrent updates to policy labels for a policy lead to an infinite loop causing the Ranger database to crash. Concurrent policy label update and underlying policy deletion also result in the same behavior.

Fixed concurrent updates to policy labels. Policy label updates are skipped if the underlying policy gets deleted.

Cloudera Runtime 7.3.1.200 SP1

CDPD-79911: Upgrade netty to 4.1.118.Final

Ranger: Upgraded netty to 4.1.118.Final due to CVE-2025-24970, CVE-2025-25193.

CDPD-79778: Issue in APPEND mode while writing Ranger audits to HDFS

Ranger audits to HDFS currently use APPEND mode in case of errors or exceptions encountered in writing audits to HDFS destination (to prevent large number of audit files), and falls back to WRITE mode if unable to APPEND.

The issue is fixed now. A configuration parameter to enable APPEND mode while writing Ranger audits to HDFS is added.

CDPD-79460: RMS full-sync breaks due to unsupported schema

RMS supports HDFS and Ozone file-system in private cloud and S3 in public cloud. The supported file schema types are hdfs, s3a, o3fs, and ofs. If Hive table location was stored at other file-system which was not supported by the RMS, then full-sync threw exception, due to unsupported schema. The full-sync was never completed.

This fix skips the unsupported file schema types while processing table and database metadata during full-sync and delta-sync in RMS. Therefore, table and database locations stored at other file-system, which is not supported by the RMS, are not mapped and appropriate messages are logged in RMS server log file.

CDPD-78195: Enhance the audit generated in Ranger during data discovery call from REST Catalog API

The audit generated in Ranger during data discovery call from the REST Catalog API is now enhanced. Calls such as list Databases / ListTables did not have the correct access Types and are enhanced to provide details on the operation.

CDPD-77093: HBase scan operation returns denied columns in result

In some cases, Ranger authorization returned access results of some HBase data even when the user was not entitled to do so.

This issue is fixed now.

Cloudera Runtime 7.3.1.100 CHF 1

CDPD-78204: Alter Rename should not check for the CREATE permission on the database in which renamed table is created

Alter rename command does not require CREATE permission on the database in which renamed table is created.

CDPD-78072: Set role command was not audited by Ranger

The issue is fixed now. Ranger now supports the auditing of the SET ROLE command for Hive plugin.

CDPD-77948: CSV injection vulnerability during CSV and Excel file export

When policies are created with the special characters mentioned in the document, there were vulnerabilities which can be exploited.

The issue is fixed now. Checks have been added to ensure whenever such characters are present, a space after it is added.

CDPD-77093: Hbase scan operation returns denied columns in result

In some cases, Ranger authorization returns access results of some HBase data even when the user is not entitled to.

This issue is fixed now.

CDPD-76662: 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-76630: Ranger Audit Filter for the HBase service was not working as expected

On the service creation page, while adding audit filters, the resources column includes an Include/Exclude switch for most resources. The issue was arising when selecting an option in the switch:

- If Include was selected, the isExclude parameter should be false, but it was incorrectly set to true.
- If Exclude was selected, the isExclude parameter should be true, but it was incorrectly set to false.

Due to this incorrect mapping, the switch values were reversed, causing the audit filter values to be set incorrectly. As a result, incorrect audit access logs were generated.

The issue is fixed now.

CDPD-76131: A ResourceTrie node referring to modified policy-evaluator was removed even when it contained wildcard-evaluator(s)

If the policy-deltas were enabled, then when two policies had a common subset of resources and were defined on same user (or subset of users, through groups or direct users), that time if one of these policies was modified (on anything: name, resource, user), it was the only one in effect during access evaluation, until the underlying service was restarted. The underlying cause was a ResourceTrie node referring to modified policy-evaluator was removed even when it contained wildcard-evaluator(s).

This fix removes self node from the resourceTrie only if it has no children, no evaluators, and no wildcard-evaluators.

CDPD-75947: Support SASL bind for Ranger Usersync with AD/LDAP

Usersync of Ranger supports GSSAPI SASL Bind. For more information see, [Support for SASL Bind in Ranger Usersync](#).

CDPD-75105: Performance fixes for Ozone plugin

Fixed the performance issues observed while evaluating policies for multi-level resources:

- RANGER-4893: Improves policy evaluation for multilevel resource hierarchies.
- RANGER-4922: Reduces time to find tags associated with multilevel resources.

CDPD-72979: Ranger Tagsync did not support Ozone OFS paths/O3FS recursive feature

There was no support for OFS path/O3FS recursive feature in 7.3.1. So while you upgraded from 7.1.9 SP1 CHF3 or higher to 7.3.1, you saw a regression.

This issue has been fixed now in 7.3.1 CHF1. Ozone keys will now be recursively checked for tags and tag based policies. So, tags applied for parent directory will be applicable to subdirectories too. If you are already using tag based policies for Ozone keys and upgrading from 7.1.9 SP1 CHF2 or lower or 7.3.1, and you want the new behavior (i.e. isRecursive=true) for old tagged keys, you need to retag these keys in Atlas.

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

[CVE-2024-55532](#) - Apache Ranger

Cloudera Runtime 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 Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

CDPD-83428: Logic to check if metric collection is thread-safe or not, not to parse the entire configuration every-time

Previously, the KMSWenApp metric collection logic tried getting and parsing the Key Management Server (KMS) configuration every time KMS configuration was required. The underlying configuration class acquired a lock for initialization blocking other threads from using the KMS configuration.

This issue is now fixed. The configuration value is now initialized and kept as an instance member, which is used every time to check metrics collection thread safety flags.

Cloudera Runtime 7.3.1.400 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.200 SP1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.100 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1**OPSAPS-70657: KEYTRUSTEE_SERVER & RANGER_KMS_KTS migration to RANGER_KMS from Cloudera Runtime 7.1.x to UCL**

KEYTRUSTEE_SERVER and RANGER_KMS_KTS services are not supported starting from the Cloudera Runtime 7.3.1.0 release. Therefore added validation and confirmation messages to the Cloudera Manager upgrade wizard to alert the user to migrate KEYTRUSTEE_SERVER keys to RANGER_KMS before upgrading to Cloudera Runtime 7.3.1.0 release.

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

The Keytrustee components - KEYTRUSTEE_SERVER and RANGER_KMS_KTS services are not supported starting from the Cloudera Runtime 7.3.1.0 release. These services cannot be installed or managed with Cloudera Manager 7.13.1.0 using Cloudera Runtime 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, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.400 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.200 SP1**CDPD-79108: The KerberosBasicAuthenticationHandler used by Schema Registry and Streams Messaging Manager can not handle colon in password**

7.3.1.200

The KerberosBasicAuthenticationHandler is now capable of handling colons in passwords.

Cloudera Runtime 7.3.1.100 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1**OPSAPS-68708: Schema Registry might fail to start if a load balancer address is specified in Ranger**

7.3.1

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

7.3.1

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`

7.3.1

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

CDPD-58949: Schemas are deduplicated during import

7.3.1

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

7.3.1

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, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3**CDPD-87276: Upgraded the poi-ooxml version to 5.4.0+**

7.3.1.500

The poi-ooxml Java library version is upgraded to 5.4.0 due to [CVE-2025-31672](#).

CDPD-83454: Solr authentication bypass issue

7.3.1.500

This fix handles a Solr authentication bypass flaw that allowed unauthorized users to access protected APIs as part of [CVE-2024-45216](#).

Apache Jira: [SOLR-17417](#)

Cloudera Runtime 7.3.1.400 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.200 SP1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.100 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1

OPSAPS-71690: Update control group V2 configuration parameters

7.3.1

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](#).

Technical Service Bulletins

TSB-847: CVE-2025-30065 Apache Parquet vulnerability

On April 1, 2025, a critical vulnerability in the parquet-avro module of Apache Parquet (CVE-2025-30065, CVSS score 10.0) was announced.

Remediation for affected versions

The Cloudera Search release patched through the Cloudera platform updates for Cloudera on cloud and Cloudera on premises.



Note: Cloudera will not provide remediation options for unsupported versions, and has not tested mitigations on unsupported versions. Customers are advised to upgrade to a supported product version. For more information, refer to the [Support Lifecycle Policy](#) page.

Vulnerability details

Exploiting this vulnerability is only possible by modifying the accepted schema used for translating Parquet files and subsequently submitting a specifically crafted malicious file.

Schema parsing in the parquet-avro module of Apache Parquet 1.15.0 and previous versions allows bad actors to execute arbitrary code. Attackers may be able to modify unexpected objects or data that was assumed to be safe from modification. Deserialized data or code could be modified without using the provided accessor functions, or unexpected functions could be invoked.

Deserialization vulnerabilities most commonly lead to undefined behavior, such as memory modification or remote code execution.

Action required - Mitigation for affected Cloudera products:

Until the upgrade with Apache Parquet 1.15.1 or higher is available:

1. Utilize a File Integrity Monitoring (FIM) solution. This allows administrators to monitor files at the filesystem level and receive alerts on any unexpected or suspicious activity in the schema configuration.
2. Monitor network activity for any transmission of Parquet files, and alert on any unexpected activity.
3. Be cautious with Parquet files from unknown or untrusted sources. If possible, do not process files with uncertain origin or that came from outside the organization.
4. Ensure that only authorized users have access to endpoints that ingest Parquet files.

For the latest update on this issue see the corresponding Knowledge Article: [TSB 2025-847: Critical Apache Parquet vulnerability CVE-2025-30065](#)

Fixed Issues in Apache Spark

This section lists the issues in Apache Spark that are fixed in Cloudera Runtime 7.3.1 release, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

There are no new fixed issues in this release.

Cloudera Runtime 7.3.1.400 SP2

CDPD-75091: Backport SPARK-47217 and related changes

7.3.1.400 SP2

Backports upstream Apache Spark improvements to enable reading Parquet files with mixed or widened types without precision loss or failures.

Apache Jira: [SPARK-47217](#)

Cloudera Runtime 7.3.1.300 SP1 CHF 1

CDPD-79763: Fix clobbering of files across epochs in Spark Structured streaming with Iceberg

7.3.1.300 SP1 CHF1

Backporting an upstream fix for a bug in structured streaming that resulted in clobbering of files in Iceberg tables by.

Cloudera Runtime 7.3.1.200 SP1

CDPD-79251: Spark - Timestamp read/write performance degradation

7.3.1.200 SP1

Fixing an issue where conversion between Spark's internal timestamp representation and Hive's Timestamp representation were slower on Spark 3 than on Spark 2.

CDPD-76849: Backport SPARK-40876 and related changes

7.3.1.200 SP1

Backporting SPARK-41096, SPARK-46092, SPARK-45604, SPARK-46466, SPARK-40876, and SPARK-48603

Apache Jira: [SPARK-41096](#), [SPARK-46092](#), [SPARK-45604](#), [SPARK-46466](#), [SPARK-40876](#), [SPARK-48603](#)

CDPD-70233: Rebase CDP 7.3.x Spark3 on Apache Spark 3.5.4

7.3.1.200 SP1

Upgrading Spark from 3.4.1 to 3.5.4. For more information, refer to [Migrating Spark applications](#).

Cloudera Runtime 7.3.1.100 CHF 1

CDPD-76229: Optimize the processing speed of BinaryArithmetic#dataType when processing multi-column data

7.3.1.100 CHF1

Restoring performance of some queries in Spark 3.4.1 to match other versions (3.3.x, 3.5.x) of Spark.

Optimized the processing speed of BinaryArithmetic#dataType when processing multi-column data.

Apache Jira: [SPARK-45071](#)

CDPD-75926: Backport SPARK-44653

7.3.1.100 CHF1

Backported SPARK-44653 to fix cache breaking with non-trivial DataFrame unions.

Apache Jira: [SPARK-44653](#)

CDPD-75755: [ENCODER_NOT_FOUND] Not found an encoder of the type T to Spark SQL internal representation when using Parameterized Bean

7.3.1.100 CHF1

Fixed an upstream regression causing Encoder Exception for a parameterized class

Apache Jira: [SPARK-46679](#)

CDPD-75622: Backport upstream fixes for handling nested beans and generic type beans while creating Spark encoders.

7.3.1.100 CHF1

Backporting upstream fixes from Spark 3.4 to fix the following issues:

- Starting from Spark 3.4.x, Encoders.bean raised an exception when the passed class contains a field whose type is a nested bean with type arguments
- From Spark 3.4.x, an exception is raised when Encoders.bean is called providing a bean having read-only properties
- Unsupported feature of bean encoder when the superclass of the bean has generic type arguments

Apache Jira: [APACHE-44634](#), [APACHE-45081](#), [APACHE-44910](#)

CDPD-75353: CHAR and VARCHAR handling in Spark 3 is incompatible with Spark 2

7.3.1.100 CHF1

Adding a new configuration `spark.cloudera.legacy.charVarcharLegacyPadding` (by default set to false in Spark 3). When set to true (together with `spark.sql.legacy.charVarcharAsString=true`) it creates compatibility with Spark 2 behavior.

For more information refer to [Migrating Spark applications](#).

CDPD-75286: Spark History UI - StreamConstraintsException: String length exceeds the maximum length

7.3.1.100 CHF1

Fixing an issue with Jackson to allow unlimited json string length in Spark event logs.

CDPD-59617: Spark - Upgrade Okio to 1.17.6 due to CVE-2023-3635

7.3.1.100 CHF1

Updating okio from version 1.15.0 to 1.17.6 to address the security vulnerability CVE-2023-3635.

CDPD-74730: Backport SPARK-46239: Hide the Jetty server's version

7.3.1.100 CHF1

The Jetty server's version is now hidden.

Apache Jira: [SPARK-46239](#)

CDPD-73233: Encoder not found of the type T to Spark SQL internal representation

7.3.1.100 CHF1

Fixing an upstream regression of encoder exception (`org.apache.spark.SparkUnsupportedOperationException: [ENCODER_NOT_FOUND]`) for generic types.

Apache Jira: [SPARK-49789](#)

Cloudera Runtime 7.3.1**CDPD-74697 - Spark Iceberg vectorized Parquet read of decimal column is incorrect**

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

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

7.3.1

Fixed Issues in Spark Atlas Connector

This section lists the issues in Spark Atlas Connector that are fixed in Cloudera Runtime 7.3.1 release, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

There are no new fixed issues in this release.

Cloudera Runtime 7.3.1.400 SP2**CDPD-73828: SAC - spark_process entity contains huge strings in its fields**

7.3.1.400 SP2

Adds a new configuration property in Cloudera Manager to allow disabling the fields details and sparkPlanDescription in the spark_process entity, which could cause Atlas OOM errors if they contained too large strings.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no new fixed issues in this release.

Cloudera Runtime 7.3.1.200 SP1

There are no new fixed issues in this release.

Cloudera Runtime 7.3.1.100 CHF 1

There are no new fixed issues in this release.

Cloudera Runtime 7.3.1

There are no new fixed issues in this release.

Fixed Issues in Apache Sqoop

Review the list of Sqoop issues that are resolved in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3:

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.400 SP2:

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.200 SP1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.100 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1

There are no fixed issues in this release.

Fixed Issues in Streams Messaging Manager

Review the list of Streams Messaging Manager issues that are resolved in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.400 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.200 SP1**CDPD-78694: Streams Messaging Manager UI cannot show replication details**

7.3.1.200

The Kafka cluster replication details view is now available on the Cluster Replications page.

CDPD-79108: The KerberosBasicAuthenticationHandler used by Schema Registry and Streams Messaging Manager can not handle colon in password

7.3.1.200

The KerberosBasicAuthenticationHandler is now capable of handling colons in passwords.

Cloudera Runtime 7.3.1.100 CHF 1**CDPD-73733: Schema version selection does not work in Data Explorer**

7.3.1.100

You can now switch between different Avro schema versions using the drop-down list in Data Explorer.

Cloudera Runtime 7.3.1**OPSAPS-71258: Kafka, Streams Replication Manager, and Streams Messaging Manager cannot process messages compressed with Zstd or Snappy if /tmp is mounted as noexec**

7.3.1

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 Streams Messaging Manager

7.3.1

Streams Messaging Manager 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

7.3.1

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

7.3.1

KNOX now sends the Authentication header on FIPS clusters.

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

7.3.1

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

7.3.1

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 issues that are resolved in Cloudera Runtime 7.3.1, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.400 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.200 SP1**CDPD-78694: Streams Messaging Manager UI cannot show replication details**

7.3.1.200

The Kafka cluster replication details view is now available on the Cluster Replications page.

Cloudera Runtime 7.3.1.100 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1

OPSAPS-71258: Kafka, Streams Replication Manager, and Streams Messaging Manager cannot process messages compressed with Zstd or Snappy if /tmp is mounted as noexec

7.3.1

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

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, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.400 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.200 SP1

COMPX-18663: Cgroup v2 support must fall back to v1 when there are no v2 controllers

7.3.1.200

Cgroup v1/v2 mixed mode support was introduced previously. If the cgroup v2 support was enabled (using `yarn.nodemanager.linux-container-executor.cgroups.v2.enabled` set to true), with only cgroup v1 controllers mounted on the Node Managers, the cgroup v2 support did not fall back to v1. This issue is now resolved and cgroup v2 support now falls back to v1 when there are no v2 controllers.

Apache Jira: [YARN-11743](#)

COMPX-18909: NodeManager marked as unhealthy if an application is terminated

7.3.1.200

Node Managers are marked unhealthy if the container-executor has an unrecoverable/configuration error, that is, the container-executor script is missing. However, if the application was terminated just before one of the containers was trying to access the localizer syslog file, causing an `IOException` then the Node Managers was marked unhealthy incorrectly. This issue is now resolved and the error checking is more specific, thereby, removing the false positive Node Managers markings.

Apache Jira: [YARN-11753](#)

COMPX-18545: Setting maximum-application-lifetime using AQCv2 templates does not apply on the first submitted application

7.3.1.200

Setting the maximum-application-lifetime property using the AQC v2 templates did not apply to the first submitted application but was applied to the subsequent ones. This issue is now resolved.

Apache Jira: [YARN-11708](#)

COMPX-18589: YARN ResourceManager raised an exception during comparison of queues

7.3.1.200

YARN ResourceManager raised an exception, java.lang.IllegalArgumentException: Comparison method violates its general contract!. The RCA was with the AND condition that caused the exception of TimSort algorithm during comparison of queues. This issue is now resolved.

Apache Jira: [YARN-11745](#)

CDPD-49702: NodeManager must be shut down when the program /var/lib/yarn-ce/bin/container-executor cannot be run

7.3.1.200

Previously, a job failed when NodeManager encountered the No such file or directory error when running the /var/lib/yarn-ce/bin/container-executor program. This issue is now resolved and NodeManager is marked as unhealthy and shut down when it cannot run the program.

Apache Jira: [YARN-11709](#)

Cloudera Runtime 7.3.1.100 CHF 1**Fixed the order of updating CPU controls with cgroup v1**

7.3.1.100

This fix ensures that cpu.cfs_period_us is updated before cpu.cfs_quota_us, to keep the ratio between the two values and not to overcome the limit defined at parent level.

Apache Jira: [YARN-11733](#)

Cloudera Runtime 7.3.1**COMPX-17702: Backport - YARN-10345 - HsWebServices containerlogs does not honor ACLs for completed jobs**

7.3.1

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

7.3.1

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 yarnuiv2 page

7.3.1

Previously, the RM did not open the Yarn application manager node web link on the **yarnuiv2** 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, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.400 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.200 SP1

There are no fixed issues in this release.

Cloudera Runtime 7.3.1.100 CHF 1

There are no fixed issues in this release.

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 Platform

Learn about the known issues in the Platform for Cloudera Runtime 7.3.1 release, the impact or changes to the functionality, and the workaround.

Known Issue in Cloudera Runtime 7.3.1

CDPD-80802, CDPD-79725: Runtime services on a running Cloudera Data Hub cluster might stop because of insufficient memory error

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



Note: You are likely to encounter this issue with the Hive service going to an error state.

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

To mitigate this issue, you can perform either of the following actions, or both:

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

Known Issues in Apache Atlas

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

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2

CDPD-83455: Basic Search API responding with unexpected values in response JSON

7.3.1.400, 7.3.1.500

The `/v2/search/basic` API endpoints can returns an unexpected result when Cloudera Data Catalog sends a search request with the equals keyword. Apache Atlas does not support the equals keyword and replies it

Only use the supported keywords when querying the `/v2/search/basic` endpoint:

- `=`
- `!=`
- `contains`
- `begin_with`
- `ends_with`
- `is_null`
- `not_null`

Apache JIRA: [ATLAS-5053](#)

Known Issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF 1

The following section lists the known issues identified in this release:

CDPD-77434: When complete qualifiedName is provided as prefixString, sometimes search suggestions do not return the right suggestion

7.1.7 and its service packs, 7.1.9 and its service packs, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

When searching for a newly created Iceberg table with the complete qualifiedName of an iceberg_ column, the quick search does not list the Iceberg table as a suggestion sometimes.

None

Known Issues identified in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF 1

The following section lists the known issues identified in this release:

CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

7.1.7. SP3 CHF7, 7.1.7. SP3 CHF8, 7.1.7. SP3 CHF9, 7.1.7. SP3 CHF10, 7.1.9 SP1 CHF4, 7.1.9 SP1 CHF5, 7.1.9 SP1 CHF6, 7.1.9 SP1 CHF7, 7.1.9 SP1 CHF8, Cloudera Runtime 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-81092: Classic UI: Tables tab view stuck on loader in entity detail page

7.3.1.300

Cloudera Runtime 7.3.1.100, 7.3.1.200

When navigating to the **Tables** tab in the details of an entity, the **Tables** tab might get stuck during loading.

None

CDPD-81499: Basic search is lagging as some Datatypes not getting indexed in Solr

7.1.9 SP1 and its service packs Cloudera Runtime 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

The Boolean datatype in Solr 3.6.2 is not indexed. When a search includes Boolean datatype, in-memory searching must be used. A larger dataset results in longer search time.

None

CDPD-82054: UI: when server response date fields as '0', UI shows as current time

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300

7.3.1.400

If an API response contains an invalid date value (such as 0) intended for display on the user interface, the current system date is shown. This issue specifically affects the **Entity Detail** page, where the create time and modified time are displayed.

Apache JIRA: [ATLAS-5015](#)

None

CDPD-84502: Advanced Search is not working properly in Atlas UI

7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400

7.3.1.500

The Advance Search in Atlas does not return any results when selecting the Typedef option. The API is not triggered for each relation.

None.

Known Issues in Cloudera Runtime 7.3.1

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

7.1.9 SP1 and its hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.1.9 SP1, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

Adding Japanese labels or user defined properties results in the error message: “Invalid label: ###, label should contain alphanumeric characters, _ or -”

None

CDPD-69279: Quick search does not return entities when using Japanese or Chinese characters to search properties

7.1.9 SP1 and its hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.1.9 SP1 and its hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.1.9 SP1 and its hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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.

7.1.7 and its service packs and hotfixes, 7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

None.

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

7.1.9 SP1 and its hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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-43772: Performance issues with Atlas service

7.1.7 and its service packs and hotfixes, 7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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-70450: Impala SQL queries that include the “WITH” clause should populate lineage in Atlas

7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200

7.3.1.300

Impala SQL queries that do not use the WITH clause can show lineage in Atlas, but queries that do use the WITH clause cannot show lineage in Apache Atlas. Impala SQL queries using the WITH clause are not supported.

CDPD-56590: Create table "like" from Iceberg table creates a hive_table instead of iceberg_table

7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

By default, for tables created using the "like" command, lineage is not generated in Atlas. The destination like table should be of the same type as source table. Instead an iceberg_table for source and hive_table for destination are getting created.

CDPD-56085: [Impala Iceberg] LOAD DATA INPATH to Iceberg_table creates a temporary hive_table with name <iceberg_table_name>_tmp* and then marks it as DELETED in Atlas

7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200

Running a query like "LOAD DATA INPATH to iceberg_table", creates a temporary hive_table with name <iceberg_table_name>_tmp* and then marks it as DELETED in Atlas. So in Atlas, a deleted entity is created corresponding to the temporary table "<iceberg_table_name>_tmp*".

Tag added to the File system (HDFS) entity will not be propagated to the Iceberg table, user has to manually add to the iceberg_table, since the tag propagation is broken due to the deleted table in the flow.

CDPD-55301: The ddlQueries and ALTERNATE_* lineage are missing for Spark tables created using spark3-shell

7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

The ddlQueries and outputFromProcesses (lineage) is missing for the alter queries.

CDPD-11941: Table creation events missed when multiple tables are created in the same Hive command

7.1.7 and its service packs and hotfixes, 7.1.8 and its service packs and hotfixes, 7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

When multiple Hive tables are created in the same database in a single command, the Atlas audit log for the database may not capture all the table creation events. When there is a delay between creation commands, audits are created as expected.

None

CDPD-11940: Database audit record misses table delete

7.1.7 and its service packs and hotfixes, 7.1.8 and its service packs and hotfixes, 7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

When a hive_table entity is created, the Atlas audit list for the parent database includes an update audit. However, at this time, the database does not show an audit when the table is deleted.

None

CDPD-11692: Navigator table creation time not converted to Atlas

7.1.7 and its service packs and hotfixes, 7.1.8 and its service packs and hotfixes, 7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400

In converting content from Navigator to Atlas, the create time for Hive tables is not moved to Atlas.

None

CDPD-10576: Deleted Business Metadata attributes appear in Search Suggestions

7.1.7 and its service packs and hotfixes, 7.1.8 and its service packs and hotfixes, 7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

Atlas search suggestions continue to show Business Metadata attributes even if the attributes have been deleted.

None

CDPD-10574: Suggestion order doesn't match search weights

7.1.7 and its service packs and hotfixes, 7.1.8 and its service packs and hotfixes, 7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

At this time, the order of search suggestions does not honor the search weight for attributes.

None

CDPD-7982: HBase bridge stops at HBase table with deleted column family

7.1.7 and its service packs and hotfixes, 7.1.8 and its service packs and hotfixes, 7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400

Bridge importing metadata from HBase fails when it encounters an HBase table for which a column family was previously dropped. The error indicates:

```
Metadata service API org.apache.atlas.AtlasClientV2$API_V2@58112bc4 failed with status 404 (Not Found) Response Body
({""errorCode"":""ATLAS-404-00-007"", ""errorMessage"":""Invalid
instance creation/updation parameters passed :
hbase_column_family.table: mandatory attribute value missing in
type hbase_column_family""})
```

None

CDPD-6675: Irregular qualifiedName format for Azure storage

The qualifiedName for hdfs_path entities created from Azure blob locations (ABFS) doesn't have the clusterName appended to it as do hdfs_path entities in other location types.

None

CDPD-4762: Spark metadata order may affect lineage

Atlas may record unexpected lineage relationships when metadata collection from the Spark Atlas Connector occurs out of sequence from metadata collection from HMS. For example, if an ALTER TABLE operation in Spark changing a table name and is reported to Atlas before HMS has processed the change, Atlas may not show the correct lineage relationships to the altered table.

None

CDPD-4545: Searches for Qualified Names with "@" doesn't fetch the correct results

When searching Atlas qualifiedName values that include an "@" character (@), Atlas does not return the expected results or generate appropriate search suggestions.

Consider leaving out the portion of the search string that includes the @ sign, using the wildcard character * instead.

CDPD-3208: Table alias values are not found in search

When table names are changed, Atlas keeps the old name of the table in a list of aliases. These values are not included in the search index in this release, so after a table name is changed, searching on the old table name will not return the entity for the table.

None

CDPD-3160: Hive lineage missing for INSERT OVERWRITE queries

Lineage is not generated for Hive INSERT OVERWRITE queries on partitioned tables. Lineage is generated as expected for CTAS queries from partitioned tables.

None

CDPD-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 access to Atlas after logging out, close all browser windows and exit the browser.

CDPD-1892: Ranking of top results in free-text search not intuitive

The Free-text search feature ranks results based on which attributes match the search criteria. The attribute ranking is evolving and therefore the choice of top results may not be intuitive in this release.

If you don't find what you need in the top 5 results, use the full results or refine the search.

CDPD-1884: Free text search in Atlas is case sensitive

The free text search bar in the top of the screen allows you to search across entity types and through all text attributes for all entities. The search shows the top 5 results that match the search terms at any place in the text (*term* logic). It also shows suggestions that match the search terms that begin with the term (term* logic). However, in this release, the search results are case-sensitive.

If you don't see the results you expect, repeat the search changing the case of the search terms.

CDPD-1823: Queries with ? wildcard return unexpected results

DSL queries in Advanced Search return incorrect results when the query text includes a question mark (?) wildcard character. This problem occurs in environments where trusted proxy for Knox is enabled, which is always the case for Cloudera.

None

CDPD-1664: Guest users are redirected incorrectly

Authenticated users logging in to Atlas are redirected to the Cloudera Knox-based login page. However, if a guest user (without Atlas privileges) attempts to log in to Atlas, the user is redirected instead to the Atlas login page.

To avoid this problem, open the Atlas Dashboard in a private or incognito browser window.

CDPD-922: IsUnique relationship attribute not honored

The Atlas model includes the ability to ensure that an attribute can be set to a specific value in only one relationship entity across the cluster metadata. For example, if you wanted to add metadata tags to relationships that you wanted to make sure were unique in the system, you could design the relationship attribute with the property "IsUnique" equal true. However, in this release, the IsUnique attribute is not enforced.

None

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

7.1.7 and its service packs and hotfixes, 7.1.8 and its service packs and hotfixes, 7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

CDPD-76035: Resource lookup for Atlas service is failing

7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.1.9 SP1 CHF3, CHF4, CHF5, CHF6, CHF7, CHF8, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400

7.1.9 SP1 CHF9, , 7.3.1.500

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

7.1.9 SP1 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.1.7 SP3 and its hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400

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.

CDPD-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails

7.1.7 and its service packs and hotfixes, 7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100 7.3.1.200

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

CDPD-76789: Creating tag with name description throws java.lang.ClassCastException

7.1.7 and its service packs and hotfixes, 7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100 7.3.1.200

Creating classification with reserved names such as "name", "description", "owner", "version", "serviceType" and "options" can lead to HTTP 500 error in Apache Atlas.

Avoid creating classification with reserved names such as "name", "description", "owner", "version", "serviceType" and "options".

CDPD-77767: Migration status is not updated on file while migration is in progress

7.3.1

7.3.1.100

Whenever a migration starts on any environment and the vertex is attempted to be updated with details, the status update fails with a null pointer exception.

None.

CDPD-71411: Atlas Hbase import fails with jackson-databind exception

7.1.9 SP1 and its service packs and hotfixes, 7.3.1

7.1.9 SP CHF4, 7.3.1.100

The jackson-annotations is not updated to match jackson-databind (2.15.0). This causes the NoSuchFieldError error intermittently during an HBase import because of a version mismatch (older 2.12.7 on classpath).

CDPD-80922: Without permission for one glossary, /glossary call throws exception

7.1.7 SP3 and its hotfixes, 7.1.9 SP1 hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200

7.3.1.300

When you don't have the permission for a glossary, the /glossary call results in an exception, blocking authorization for all glossaries. The error message is the following:

```
{  "errorCode": "ATLAS-403-00-001",  "errorMessage": "hrt_qa is
not authorized to perform read entity: guid=*****_****_****_
****_*****"}

```

CDPD-69910: NPE while deleting BusinessMetadata

7.1.7 and its service packs and hotfixes, 7.1.9 and its service packs and hotfixes, 7.3.1

7.3.1.100

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

None

Apache Jira: [ATLAS-4863](#)

CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export

7.1.9 and its service packs and hotfixes, 7.3.1 and its service packs and hotfixes

When creating tables with multiple levels of depth, a tag applied to the first table will be propagated along the lineage. After exporting the last table, it will not have the propagated tag.

Apache JIRA: [ATLAS-4889](#)

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

7.1.9 and its service packs and hotfixes, 7.3.1 and its service packs and hotfixes

When creating tables with multiple levels of depth, a tag applied to the first table will be propagated along the lineage. After dropping the first table and exporting the entire lineage, the child tables will not have the propagated tag.

Apache JIRA: [ATLAS-5055](#)

DOCS-26821: Iceberg entities created in 7.2.17 Atlas-Hook (Hive,Impala,Spark) is not ingested by 7.2.18/7.3.1 Atlas Server

When using an Atlas hook from 7.3.1 with an earlier version of Atlas server (7.1.9 SP1 CHF5 or earlier) to create an Iceberg table, the older data model of Iceberg cannot process the features of the new data models. For example, when encountering relationships or attributes in queries which do not have a definition in the earlier model, Iceberg fails to create new tables with an error message. The data model of the Iceberg in Cloudera Runtime 7.1.9 SP1 CHF5 or later is not backward compatible with earlier data models.

It is recommended to upgrade your Atlas server version to at least 7.1.9 SP1 CHF5 or later.

CDPD-87670: [Atlas UI] Apache Atlas Glossary becomes unresponsive when the page size is set to 50

7.1.7 and its service packs and hotfixes, 7.1.9 and its service packs and hotfixes, 7.3.1 and its service packs and hotfixes

When the page limit is set to 50 in **Glossary**, the user interface becomes unresponsive due to an incorrect API call with the error message, when trying to show more than 25 associated entities for a glossary item: "expected type `AtlasGlossaryCategory`; found `AtlasGlossaryTerm`".

Apache Jira: [ATLAS-5067](#)

CDPD-79099: Unable to import hive tables using import script in Public Cloud

7.1.7 and its service packs and hotfixes, 7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400

7.3.1.500

The validation of import .zip files can cause the `org.apache.atlas.AtlasServiceException` error if the .zip file contains Atlas shell entities missing mandatory attributes such as a name. Therefore, the import of shell entities fails.

CDPD-77128: Log out message for Atlas needs to be more informative on Knox enabled cluster

7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500
7.3.1.600

Logging out of Atlas does not manage the external authentication. Atlas does not communicate a log-out event with the external authentication management, Apache Knox. When you log out of Atlas, you can still open the instance of Atlas from the same web browser without re-authentication.

To prevent additional access to Atlas, close all browser windows and exit the browser.

CDPD-83398: REMOVE_LEGACY_REF_ATTRIBUTES Patch Skipped During Startup, Causing Errors While Processing DML Events

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

The REMOVE_LEGACY_REF_ATTRIBUTES patch is skipped during Atlas patch processing, which leads to an error when Data Manipulation Language (DML) events (e.g., "insert data") are processed, events that are typically ignored. This is observed in environments where DML event filtering is not in place or the patch is not applied correctly. As a result, Atlas attempts to process these events and encounters errors related to missing or outdated attribute handling.

CDPD-88230: Bulk import operations currently do not construct lineage as expected

7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

Bulk import operations do not create lineage consistently, even when the entities are imported successfully.

CDPD-88302: Bulk import operations currently do not create audits for the imported entities

7.1.9 and its service packs and hotfixes, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

Bulk import operations do not create audits for the imported entities.

OPSAPS-74147: ZDU | Upgrade from 7.1.7.3000 to 7.3.1.400 fails

7.1.7 and its service packs and hotfixes, 7.3.1.400

The ZDU rolling upgrade for Atlas fails from Cloudera Runtime 7.1.7 SP3 to 7.3.1.x version, as the Zero Downtime Upgrade feature is not available for 7.1.7 SP3. The RoleState for Atlas is not checked and the Update Solr Configs for Atlas command fails.

Known Issues in Apache Avro

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

Known issues identified in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.100 CHF1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1

There are no new known issues identified in this release.

Known Issues in Cloud Connectors

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

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF 1

There are no new known issues identified in this release.

Cloudera Runtime 7.3.1

AWS SDK 2.25.53 warning about transfer manager not using CRT client

7.3.1.0

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 S3AsyncC
lient 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

CDPD-76378 Uploading files to S3 takes longer than expected

7.1.9 SP 1, 7.3.1.0

7.3.1.100

Uploading files to S3 on long haul networks is delayed if the file is smaller than 64MB. As the 100-continue header is not received immediately, the files are uploaded only after several retries. The delay in uploading is caused by the `fs.s3a.connection.expect.continue` S3A configuration parameter. The `fs.s3a.connection.expect.continue` S3A configuration property controls whether or not a PUT request to the S3 store sets the "Expect: 100-continue" header, and awaits a 100 CONTINUE response before uploading any data. This allows throttling and other transient problems to be reported by the S3 store before there is an attempt by the client to upload any data.

By default, the value of the configuration property is set to true, which means that the header is set, and a 100 CONTINUE response is required.

If there is a significant delay in uploading files or if the write time of jobs increases, you can set the `fs.s3a.connection.expect.continue` configuration property for your service to `false`.

Known issues in Cruise Control

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

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF 1

There are no new known issues identified in this release.

Cloudera Runtime 7.3.1

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

7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

If the CPU usage metric is not reported, the `numValidWindows` in Cruise Control will be 0 and proposal generation as well as partition rebalancing will not work. If this issue is present, the following message will be included in the Kafka logs:

```
WARN com.linkedin.kafka.cruisecontrol.metricsreporter.CruiseControlMetricsReporter:
    [CruiseControlMetricsReporterRunner]: Failed reporting CPU
    util.
```

```
java.io.IOException: Java Virtual Machine recent CPU usage is not
    available.
```

This issue is only known to affect Kafka broker hosts that have the following specifications:

- CPU: Intel(R) Xeon(R) CPU E5-2699 v4 @ 2.20GHz
- OS: Linux 4.18.5-1.el7.elrepo.x86_64 #1 SMP Fri Aug 24 11:35:05 EDT 2018 x86_64
- Java version: 8-18

Move the broker to a different machine where the CPU is different. This can be done by moving the host to a different cluster. For more information, see [Moving a Host Between Clusters](#)



Note: Cluster nodes affected by this issue are not displayed as unhealthy.

Known Issues in Apache HBase

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

Known issues identified in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1.100 CHF 1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1

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

7.3.1 and its higher versions

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.

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

7.3.1 and its higher versions

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

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

7.3.1 and its higher versions

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.

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

HDFS encryption with HBase

7.3.1 and its higher versions

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.

N/A

Snappy compression with `/tmp` directory mounted with `noexec` option

7.3.1 and its higher versions

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

7.3.1 and its higher versions

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 `postOperation` 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.

N/A

Apache Issue: [HBASE-6992](#)

HBase shutdown can lead to inconsistencies in META

7.3.1 and its higher versions

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.

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

7.3.1 and its higher versions

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.

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

7.3.1 and its higher versions

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

N/A

Apache Issue: N/A

CDPD-77399: HBase fails to register the servlet metrics and throws ClassNotFoundException: org.apache.hadoop.metrics.MetricsServlet

7.3.1, 7.3.1.100 CHF 1

7.3.1.200 SP1

The MetricsServlet class is a Hadoop 2-based metric servlet unavailable in Hadoop 3 deployments.

Ignore this WARN log message during HBase Master and RegionServer startup.

Apache Issue: [HBASE-28315](#)

HBase throws a NullPointerException on MemstoreFlusher when flush is triggered by too many WALs after a WAL rolling

7.3.1 and its higher versions

When rolling the current wall, if HBase reaches `hbase.regionserver.maxlogs`, a memstore flush is triggered, and HBase uses `SequenceIdAccounting.findLower` to retrieve the region and stores with the lower sequence ID in the WAL. The problem occurs when there is a flush entry in the WAL. `SequenceIdAccounting.findLower` does not filter the entries, and HBase returns a `METAFAMILY` column family, which later cannot be resolved by `HRegion.getSpecificStores`, resulting in a `NullPointerException`.

1. Increase the maximum number of WAL files.
2. Perform periodic manual flushes so that the flush is not triggered due to the *Too many WALs* condition.

Known Issues in HDFS

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

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF 1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1

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	3	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:

- 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](#))

Known Issues in Apache Hive

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

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3

CDPD-88865: Unicode character support with a MySQL backend

Cloudera on premises 7.3.1.500 and all base versions

When a cluster's backend is a MySQL database, CREATE TABLE statements with more than two Unicode column names can fail. This is a known bug, HIVE-18083, where the MySQL database does not support non-ASCII characters in column names.

None

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified for Hive in this release.

Known Issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF1

There are no new known issues identified for Hive in this release.

Known issues identified in Cloudera Runtime 7.3.1.200 SP 1

There are no new known issues identified for Hive in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF1

The following section lists the known issues identified in this release:

CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

7.1.9 SP1 CHF4, 7.1.7 SP3 CHF7, 7.3.1.100, and its higher versions

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-78490: HiveCreateSysDb command fails

7.1.9 SP1, 7.3.1.0, and its higher versions

Hive services fail to start due to HiveCreateSysDb command failure during the first run.

None

CDPD-72605: Optimizing partition authorization in HiveMetaStore

7.3.1.200

7.3.1.0

The `add_partitions()` API in HiveMetaStore unnecessarily authorizes both new and existing partitions, increasing processing time and load on the authorization service.

None

Known Issues identified in Cloudera Runtime 7.3.1

CDPD-74680: DAG not retried after failure

7.3.1 and its higher versions

When executing a Hive query, if the ApplicationMaster container fails, Hive does not retry the DAG if the failure message contains some diagnostic information including a line break, leading to query failure (instead of retry).

None

HiveServer2 goes into a hung state intermittently

7.3.1 and its higher versions

HiveServer2 can intermittently hang or crash due to heap out-of-memory (OOM) errors triggered by the default 1 GB cache limit for fetch tasks. This occurs with certain queries that exceed the available heap space.

1. Disable the fetch task caching feature by setting: `hive.fetch.task.caching=false`.
2. You can adjust the `hive.fetch.task.conversion.threshold` property to a lower value in the megabyte range. The default value in Cloudera on premises 7.3.1 is 1 GB.

Known Issues in Apache Impala

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

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3:

DWX-20490: Impala queries fail with "Caught exception The read operation timed out, type=<class 'socket.timeout'> in ExecuteStatement"

7.3.1.500

Queries in `impala-shell` fail with a socket timeout error in `execute statement` which submits the query to the coordinator. The error occurs when query execution takes longer to start mainly when query planning is slow due to frequent metadata changes.

Increase the socket timeout on the client side. Set `--client_connect_timeout_ms` to a higher value, e.g. add `--client_connect_timeout_ms=600000` to the `impala-shell` command line.

DWX-20491: Impala queries fail with EOFException: End of file reached before reading fully

7.3.1.500

Impala queries fail with an `EOFException` when reading from an HDFS file stored in an S3A location. The error occurs when the file is removed. If the file is removed using SQL commands like `DROP PARTITION`, there may be a significant lag in Hive Metastore event processing. If removed by non-SQL operations, run `REFRESH` or `INVALIDATE METADATA` on the table to resolve the issue.

Run `REFRESH/INVALIDATE METADATA <table>;`

CDPD-90807: Thrift protocol limitation during Impala zero downtime upgrade (ZDU)

7.3.1.500

Zero Downtime Upgrades (ZDU) for Impala are not supported when upgrading to version 7.3.1.500 SP3. This is due to Thrift protocol incompatibilities that can cause queries to fail during upgrade.

None

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2:

There are no new known issues identified for Impala in this release.

Known Issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF1

There are no new known issues identified for Impala in this release.

Known Issues identified in Cloudera Runtime 7.3.1.200 SP1

CDPD-80166: Ignore CREATE_TABLE events for inaccessible databases to prevent event processor error

7.3.1.200

7.3.1

CREATE_TABLE events for databases restricted by authorization caused the Impala event processor to enter an ERROR state

The event processor now ignores such events when the database is not found, logging the issue for administrators.

Apache Jira: [IMPALA-11735](#)

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF1

There are no new known issues identified for Impala in this release.

Known Issues identified in Cloudera Runtime 7.3.1

IMPALA-532: Impala should tolerate bad locale settings

7.3.1 and its higher versions

If the LC_* environment variables specify an unsupported locale, Impala does not start.

Add LC_ALL="C" to the environment settings for both the Impala daemon and the Statestore daemon.

IMPALA-691: Process mem limit does not account for the JVM's memory usage

7.3.1 and its higher versions

Some memory allocated by the JVM used internally by Impala is not counted against the memory limit for the impalad daemon.

To monitor overall memory usage, use the top command, or add the memory figures in the Impala web UI /memz tab to JVM memory usage shown on the /metrics tab.

IMPALA-635: Avro Scanner fails to parse some schemas

7.3.1 and its higher versions

The default value in Avro schema must match type of first union type, e.g. if the default value is null, then the first type in the UNION must be "null".

Swap the order of the fields in the schema specification. For example, use ["null", "string"] instead of ["string", "null"]. Note that the files written with the problematic schema must be rewritten with the new schema because Avro files have embedded schemas.

IMPALA-1024: Impala BE cannot parse Avro schema that contains a trailing semi-colon

7.3.1 and its higher versions

If an Avro table has a schema definition with a trailing semicolon, Impala encounters an error when the table is queried.

Remove trailing semicolon from the Avro schema.

IMPALA-1652: Incorrect results with basic predicate on CHAR typed column

7.3.1 and its higher versions

When comparing a CHAR column value to a string literal, the literal value is not blank-padded and so the comparison might fail when it should match.

Use the RPAD() function to blank-pad literals compared with CHAR columns to the expected length.

IMPALA-1821: Casting scenarios with invalid/inconsistent results

7.3.1 and its higher versions

Using a CAST() function to convert large literal values to smaller types, or to convert special values such as NaN or Inf, produces values not consistent with other database systems. This could lead to unexpected results from queries.

None

IMPALA-2005: A failed CTAS does not drop the table if the insert fails

7.3.1 and its higher versions

If a CREATE TABLE AS SELECT operation successfully creates the target table but an error occurs while querying the source table or copying the data, the new table is left behind rather than being dropped.

Drop the new table manually after a failed CREATE TABLE AS SELECT

IMPALA-3509: Breakpad minidumps can be very large when the thread count is high

7.3.1 and its higher versions

The size of the breakpad minidump files grows linearly with the number of threads. By default, each thread adds 8 KB to the minidump size. Minidump files could consume significant disk space when the daemons have a high number of threads.

Add `-\\-minidump_size_limit_hint_kb=size` to set a soft upper limit on the size of each minidump file. If the minidump file would exceed that limit, Impala reduces the amount of information for each thread from 8 KB to 2 KB. (Full thread information is captured for the first 20 threads, then 2 KB per thread after that.) The minidump file can still grow larger than the "hinted" size. For example, if you have 10,000 threads, the minidump file can be more than 20 MB.

IMPALA-4978: Impala requires FQDN from hostname command on Kerberized clusters

7.3.1 and its higher versions

The method Impala uses to retrieve the host name while constructing the Kerberos principal is the `gethostname()` system call. This function might not always return the fully qualified domain name, depending on the network configuration. If the daemons cannot determine the FQDN, Impala does not start on a Kerberized cluster.

Test if a host is affected by checking whether the output of the `hostname` command includes the FQDN. On hosts where `hostname` only returns the short name, pass the command-line flag `##hostname=FULLY_QUALIFIED_DOMAIN_NAME` in the startup options of all Impala-related daemons.

IMPALA-6671: Metadata operations block read-only operations on unrelated tables

7.3.1 and its higher versions

Metadata operations that change the state of a table, like `COMPUTE STATS` or `ALTER RECOVER PARTITIONS`, may delay metadata propagation of unrelated unloaded tables triggered by statements like `DESCRIBE` or `SELECT` queries.

None

IMPALA-7072: Impala does not support Heimdal Kerberos

None

CDPD-28139: Set spark.hadoop.hive.stats.autogather to false by default

As an Impala user, if you submit a query against a table containing data ingested using Spark and you are concerned about the quality of the query plan, you must run `COMPUTE STATS` against such a table in any case after an ETL operation because `numRows` created by Spark could be

incorrect. Also, use other stats computed by COMPUTE STATS, e.g., Number of Distinct Values (NDV) and NULL count for good selectivity estimates.

For example, when a user ingests data from a file into a partition of an existing table using Spark, if `spark.hadoop.hive.stats.autogather` is not set to `false` explicitly, `numRows` associated with this partition would be 0 even though there is at least one row in the file. To avoid this, the workaround is to set `"spark.hadoop.hive.stats.autogather=false"` in the "Spark Client Advanced Configuration Snippet (Safety Valve) for `spark-conf/spark-defaults.conf`" in Spark's CM Configuration section.

IMPALA-2422: % escaping does not work correctly when occurs at the end in a LIKE clause

7.3.1 and its higher versions

If the final character in the RHS argument of a LIKE operator is an escaped \% character, it does not match a % final character of the LHS argument.

None

IMPALA-2603: Crash: impala::Coordinator::ValidateCollectionSlots

A query could encounter a serious error if includes multiple nested levels of INNER JOIN clauses involving subqueries.

None

IMPALA-3094: Incorrect result due to constant evaluation in query with outer join

7.3.1 and its higher versions

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 alltypestiny a1
  INNER JOIN alltypesagg a2 ON a1.smallint_col = a2.year AND fals
e
  RIGHT JOIN alltypes a3 ON a1.year = a1.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. 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-59625: Impala shell in RHEL 9 with Python 2 as default does not work

7.1.9, 7.3.1 and its higher version

If you try to run `impala-shell` on RHEL 9 by setting the default python executable available in `PATH` to Python 2, it will fail since RHEL 9 is compatible only with Python 3.

If you run into such issues, set this parameter pointing to Python 3, `IMPALA_PYTHON_EXECUTABLE=python3`.

CDPD-42958: After upgrading the CDH 7.1.9 from CDH 6.x, under certain conditions you cannot insert data into a table

7.1.8, 7.1.9, 7.3.1 and its higher version

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.

Impala api calls via Knox require configuration if the Knox customized Kerberos principal name is a default service user name

To access Impala API calls via Knox, if the Knox customized Kerberos principal name is a default service user name, then configure "authorized_proxy_user_config" by clicking Clusters->Impala->configuration. Include the Knox customized Kerberos principal name in the comma separated list of values <knox_custom_kerberos_principal_name>="*" where <knox_custom_kerberos_principal_name> is the value of the Kerberos Principal in the Knox service. Select Clusters>Knox>Configuration and search for Kerberos Principal to display this value.

CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes.

7.1.7

You must use a single Knox node to access Impala UI.

CDPD-21828: Multiple permission assignment through grant is not working

7.1.7

None

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

7.3.1, 7.3.1.300 and its higher version

7.3.1.100, 7.3.1.200

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to cm_hdfs -> all-path for the user impala, inserting into a table whose HDFS POSIX permissions happen to exclude impala access will result in "AnalysisException: Unable to INSERT into target table (default.t1) because Impala does not have WRITE access to HDFS location: hdfs://XXXXXXXXXXXXX"

OPSAPS-46641: A single parameter exists in Cloudera Manager for specifying the Impala Daemon Load Balancer. Because BDR and Hue need to use different ports when connecting to the load balancer, it is not possible to configure the load balancer value so that BDR and Hue will work correctly in the same cluster.

The workaround is to use the load balancer configuration either without a port specification, or with the Beeswax port: this will configure BDR. To configure Hue use the "Hue Server Advanced Configuration Snippet (Safety Valve) for impalad_flags" to specify the the load balancer address with the HiveServer2 port.

Known Issues in Hue

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

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3:

CDPD-88964: Hue Logs Missing in Hue UI

7.3.1.500

You may experience an issue where the Hue UI displays only a few lines of logs instead of the complete Hue logs. This can occur due to leftover Gunicorn processes that interfere with the proper logging and display of logs within the Hue interface

1. Stop the Hue service.
2. Terminate any remaining Gunicorn processes to clear hung or orphan processes that may be causing the issue. Run the following command (use sudo if not running as root):

```
# pids=$(ps -efwww | grep rungunicornserver | grep -v "grep"
| grep "rungunicornserver"| awk '{ print $2 }') && for i in
$pids; do kill -9 $i ; done
```

3. Restart the Hue service.

CDPD-90510: Defunct Hue gunicorn worker processes

7.3.1.500

Hue on Ubuntu 22 using Oracle Database can accumulate defunct rungunicornserver worker processes due to incomplete process termination and stale database connections. This led to a cluttered process table, which does not critically impact service functionality.

Periodically clean defunct worker processes with:

```
pgrep -f 'hue rungunicornserver' | xargs -r kill -9
```

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2:

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF 1

CDPD-83015: Issues with files or directories with the % character may fail to open or Copy

7.3.1.300, 7.3.1.400, 7.3.1.500

On RHEL 9.5, files or directories containing the %character may fail to open or copy due to Apache HTTPD version 2.4.62.

None.

Known Issues identified in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF 1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1

OPSAPS-75134: LDAP and Kerberos dual authentication fails with HiveOnTez in HTTP Transport Mode

7.3.1 and its higher versions, Cloudera Manager 7.13.1 and its higher versions

Enabling LDAP for the HiveOnTez service in a Kerberos environment with transport mode set to HTTP, Hue fails to load database information due to unsupported server authentication combinations. The HiveOnTez requires `hive.server2.authentication=LDAP, KERBEROS` parameter, but Hue supports the KERBEROS and LDAP values separately, but not combined, causing a conflict and preventing a successful connection to HiveServer2.

- Log into Cloudera Manager Hive-On-Tez Configuration Hive Service Advanced Configuration Snippet (Safety Valve) for hive-site.xml and set the following value:

```
hive.server2.authentication=LDAP,KERBEROS
```

- Go to Hue Configuration Hue Server Advanced Configuration Snippet (Safety Valve) for hive-site.xml and set the following value:

```
hive.server2.authentication=KERBEROS
```

OPSAPS-69659: Hue service fails on restart with "Unable to find pycogp2 2.5.4" error

7.1.7 SP1 and its CHFs

Hue service fails to restart and you see the following error: Unable to find pycogp2 2.5.4. This could be because you have installed Python in a non-default location and Hue is unable to locate the pycogp2 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.

- Log in to Cloudera Manager as an Administrator.
- 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

- Click Save Changes.
- Restart the Hue service.

OPSAPS-73942: The upgrade fails due to configuration issues with the query processor service.

7.3.1 and its higher versions

When upgrading from CDP Private Cloud Base 7.1.x versions with the Query Processor service installed to Cloudera 7.3.1, the upgrade wizard fails during the configuration validation phase with the following warning:

"The version of the service query-processor can't be upgraded. You must remove the service before upgrading to the Cloudera 7.3.1."

To proceed with the upgrade:

- Remove the **Query Processor** service from the cluster.
- Perform the upgrade to **Cloudera 7.3.1**.
- Re-add the **Query Processor** service post-upgrade if needed.

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

7.1.9 SP1 and its CHFs, 7.3.1 and its higher versions

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

7.1.9 SP1 and its CHFs, 7.3.1 and its higher versions

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

7.1.9 SP1 and its CHF, 7.3.1 and its higher versions

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-39330: Unable to use the pip command in Cloudera

7.1.7 SP1 and its CHF

You may not be able to use the pip command in Cloudera 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 Cloudera](#).

Hue UI is blank upon login after upgrading to Cloudera 7.1.7 from CDH 6

7.1.7 SP1 and its CHF

If your cluster was secured using Knox, and if you have upgraded from CDH 6 to Cloudera 7.1.7, then you may see a blank Hue screen. This could happen because the `knox_proxyhosts` parameter is newly introduced in Cloudera, 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 Cloudera 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.

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.

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 Cloudera Base on premises.

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 Cloudera 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.

Known issues identified in Cloudera Runtime 7.3.1.500 SP3

There are no new issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.400 SP2

CDPD-83022: Incorrect row count displayed in table metadata after compaction

7.3.1.400, 7.3.1.500

After running data compaction operations on large tables, the row count displayed by the `DESC RIBE FORMATTED` command may be inaccurate. Initially, the count may appear higher than the actual number of rows. Subsequently, after running the `ANALYZE` command to update table statistics, the count might then appear lower than the actual number of rows.

This issue has been observed in large tables containing a significant number of historical snapshots (exceeding 10000). All these snapshots are primarily generated through `UPDATE` operations.

It is important to note that this is just a metadata display issue, and there is no loss of data. The underlying table data remains complete and correct.

To obtain an accurate row count, use the `SELECT COUNT(*)` query.

Known issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF1

There are no new issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.200 SP1

There are no new issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.100 CHF1

CDPD-78381: Performance degradation noticed in some Hive Iceberg TPC-DS queries

7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

While running Hive TPC-DS (Parquet + Iceberg) performance benchmarking for Cloudera Runtime 7.3.1.100, the overall performance of Iceberg tables resulted in a 15.68% increase as compared to Iceberg tables in Cloudera Runtime 7.3.1.0. However, it was noticed that some of the queries resulted in a decreased performance.

None.

CDPD-78134: CBO fails when a materialized view is dropped but its pre-compiled plan remains in the registry.

7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

Consider a cluster having two HiveServer (HS2) instances. Each HS2 instance contains its own Materialized View (MV) registry and the registries contain pre-compiled plans of MVs that are enabled for query rewriting. Without the registries, MVs will have to be loaded and compiled during each query compilation, resulting in slow query performance.

When MVs are created or dropped, they are added to or removed from the registry pertaining to the HS2 instance that issues the create or drop statement. The other HS2 instance is not immediately notified of the change. A background process is scheduled to refresh the registry, however, this process does not handle the removal of dropped MVs.

When an MV is dropped by one of the HS2 instances, it remains in the registry of the other HS2 instance. Now, if a query is processed in the second HS2 instance, the rewrite algorithm still attempts to use the dropped MV. If this MV is stored in an Iceberg table, the storage handler tries to refresh the MV metadata from the metastore but throws an exception because the MV no longer exists, resulting in a CBO failure.

Perform one of the following workarounds to address the issue:

- Restart all the HS2 instances after dropping the MV.
- From Cloudera Manager, go to `Clusters Hive Configuration` and add the `hive.server2.materializedviews.registry.impl=DUMMY` property in the HiveServer2 Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`. The DUMMY value indicates that MVs should not be cached and requests should be forwarded to Hive Metastore.



Note: Setting this property to DUMMY was done for testing purposes and can greatly increase the query compilation time.

Apache JIRA: [HIVE-28773](#)

Known issues in Cloudera Runtime 7.3.1

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

7.3.1

7.3.1.100

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

7.3.1

7.3.1

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-75411: SELECT COUNT query on an Iceberg table in AWS times out

7.3.1, 7.3.1.100, 7.3.1.200

7.3.1.300

In an AWS environment, a `SELECT COUNT` query that is run on an Iceberg table times out because some 4KB ORC file parts cannot be downloaded. This issue occurs because Iceberg uses the positional delete index only if the count of positional deletes are less than a threshold value which is by default, 100000.

None.

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

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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_B
UCKET() only takes STRING/CHAR/VARCHAR/BINARY/INT/LONG/DECIMAL/F
LOAT/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.

CDPD-84220: Cannot query Iceberg tables

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

You cannot query existing Iceberg tables after you enable HDFS HA. This is because Iceberg stores the table path in the manifest files differently depending on whether the HDFS HA is enabled or not. After you enable HDFS HA, you might not be able to query the tables created prior to you enabling HDFS HA.

None.

Known Issues in Apache Kafka

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

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF 1

There are no new known issues identified in this release.

Cloudera Runtime 7.3.1

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

7.1.7, 7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

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

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

KAFKA-2561: Performance degradation when SSL Is enabled

7.1.7, 7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

In some configuration scenarios, significant performance degradation can occur when SSL is enabled. The impact varies depending on your CPU, JVM version, Kafka configuration, and message size. Consumers are typically more affected than producers.

Configure brokers and clients with `ssl.secure.random.implementation = SHA1PRNG`. It often reduces this degradation drastically, but its effect is CPU and JVM dependent.

RANGER-3809: Idempotent Kafka producer fails to initialize due to an authorization failure

7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

Kafka producers that have idempotence enabled require the Idempotent Write permission to be set on the cluster resource in Ranger. If permission is not given, the client fails to initialize and an error similar to the following is thrown:

```
org.apache.kafka.common.KafkaException: Cannot execute transactional method because we are in an error state
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeFailWithError(TransactionManager.java:1125)
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeAddPartition(TransactionManager.java:442)
    at org.apache.kafka.clients.producer.KafkaProducer.doSend(KafkaProducer.java:1000)
    at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:914)
    at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:800)
    .
    .
    .
    Caused by: org.apache.kafka.common.errors.ClusterAuthorizationException: Cluster authorization failed.
```

Idempotence is enabled by default for clients in Kafka 3.0.1, 3.1.1, and any version after 3.1.1. This means that any client updated to 3.0.1, 3.1.1, or any version after 3.1.1 is affected by this issue.

This issue has two workarounds, do either of the following:

- Explicitly disable idempotence for the producers. This can be done by setting `enable.idempotence` to `false`.
- Update your policies in Ranger and ensure that producers have Idempotent Write permission on the cluster resource.

CDPD-49304: AvroConverter does not support composite default values

7.1.7, 7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

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

7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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: o
rg.apache.hadoop.util.NativeCodeLoader.buildSupportsSnappy()
```

```
Failed to write to HDFS due to java.lang.RuntimeException: nativ
e snappy library not available: this version of libhadoop was bu
ilt without snappy support.
```

Download and deploy missing libraries.



Important: Ensure that you complete steps 1-11 on all Kafka Connect hosts. Additionally, ensure that the advanced configuration snippet in step 12 is configured for all Kafka Connect role instances.

1. Create the /opt/nativelibs directory.

```
mkdir /opt/nativelibs
```

2. Change the owner to kafka.

```
chown kafka:kafka /opt/nativelibs
```

3. Locate the directory containing the Hadoop native libraries and copy its contents to the directory you created.

```
cp /opt/cloudera/parcels/CDH/lib/hadoop/lib/native/* /opt/nativelibs
```

4. Verify that libsnappy.so was copied to the directory you created.
5. Remove the following from /opt/nativelibs.

```
libhadoop.a
libhadoop.so
libhadoop.so.1.0.0
```

6. Run the following command.

```
hadoop version
```

The command returns the Hadoop version running in the cluster. Note down the first three digits in the version.

7. Go to <https://archive.apache.org/dist/hadoop/common/> and download the Hadoop version that matches the first three digits of the version running in the cluster.

For example, if your Hadoop version is 3.1.1.7.1.9.0-296, then you need to download Hadoop 3.1.1.

8. Extract the downloaded archive.
9. Copy the following libraries from the downloaded archive to /opt/nativelibs on the cluster host.

```
libhadoop.a          libhadoop.so.1.0.0
```

The libraries are located in `hadoop-[***VERSION***]/lib/native`.

10. Create a symlink named `libhadoop.so` and point it to `/opt/nativelibs/libhadoop.so.1.0.0`.

```
ln -s /opt/nativelibs/libhadoop.so.1.0.0 /opt/nativelibs/libhadoop.so
```

11. Change the owner of every entry within `/opt/nativelibs` to `kafka`.

```
chown -h kafka:kafka /opt/nativelibs/*
```

12. In Cloudera Manager, go to `Kafka service Configuration`.
13. Add the following key-value pair to `Kafka Connect Environment Advanced Configuration Snippet (Safety Valve)`.
 - Key: `LD_LIBRARY_PATH`
 - Value: `/opt/nativelibs`
14. Click `Save Changes`.
15. Restart the `Kafka` service.

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

7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

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

Unsupported Features

The following `Kafka` features are not supported in Cloudera

- Only `Java` and `.Net` based clients are supported. Clients developed with `C`, `C++`, `Python`, and other languages are currently not supported.
- The `Kafka` default authorizer is not supported. This includes setting `ACLs` and all related `APIs`, broker functionality, and command-line tools.
- `SASL/SCRAM` is only supported for delegation token based authentication. It is not supported as a standalone authentication mechanism.
- `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 performance to degrade

If the `Kafka` service operates with a large number of partitions, collection of partition level metrics may cause `Cloudera Manager` performance to degrade.

If you are observing performance degradation and your cluster is operating with a high number of partitions, you can choose to disable the collection of partition level metrics.



Important: If you are using Streams Messaging Manager to monitor Kafka or Cruise Control for rebalancing Kafka partitions, be aware that both Streams Messaging Manager and Cruise Control rely on partition level metrics. If partition level metric collection is disabled, Streams Messaging Manager will not be able to display information about partitions. In addition, Cruise Control will not operate properly.

Complete the following steps to turn off the collection of partition level metrics:

1. Obtain the Kafka service name.
 - a. In Cloudera Manager, Select the Kafka service.
 - b. Select any available chart, and select Open in Chart Builder from the configuration icon drop-down.
 - c. Find \$SERVICENAME= near the top of the display.
The Kafka service name is the value of \$SERVICENAME.
2. Turn off the collection of partition level metrics.
 - a. Go to Hosts Hosts Configuration .
 - b. Find and configure the Cloudera Manager Agent Monitoring Advanced Configuration Snippet (Safety Valve) configuration property.

Enter the following to turn off the collection of partition level metrics:

```
[KAFKA_SERVICE_NAME]_feature_send_broker_topic_partition_entity_update_enabled=false
```

Replace [KAFKA_SERVICE_NAME] with the service name of Kafka obtained in step 1. The service name should always be in lower case.

- c. Click Save Changes.

Known Issues in Apache Knox

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

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3

CDPD-90943: Flink Dashboard not loading correctly through Knox

7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

The Apache Flink Dashboard page does not load correctly when accessed through Knox due to a missing trailing slash (/) in the context field of the following file: /opt/cloudera/parcels/CDH/lib/knox/data/services/flink-dashboard/1.11.1/services.xml

Change /flink/dashboard to /flink/dashboard/

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2

CDPD-84236: Token generated by one Knox host fails with Unknown token error on another Knox host in Data Engineering High Availability clusters

7.3.1.400, 7.3.1.500

In Data Engineering High Availability clusters, a token generated by one Knox host may fail with an Unknown token error when accessed through another Knox host. This issue occurs due to a race condition in the PostgreSQL database, which prevents one of the Knox instances from properly initializing its configured token state service.

Restart Knox on all hosts.

Known Issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF1**OPSAPS-73038: False-positive port conflict error message appears in Cloudera Manager**

7.3.1.500

7.3.1.300, 7.3.1.400

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

None.

Known Issues identified in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified for Knox in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF1

There are no new known issues identified for Knox in this release.

Known Issues in Cloudera Runtime 7.3.1**CDPD-76294: Knox service can not be started in a large size Private Cloud Base Cloudera Manager cluster**

7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

For a large size Private Cloud Base Cloudera Manager cluster installed with Cloudera Runtime 7.3.1.0, you might face the problem that Knox can not be started with the following error message:

```
Wait Until Knox Gateway Can Serve Requests failed on Knox Gateway
```

Increase the Knox configuration parameter Knox Gateway Initial/Max Heapsize from 1 GiB to 2 GiB or 4 GiB, depending on the cluster size. Then save changes and run Restart Stale Services. After these steps, the Knox service can be started.

CDPD-71751: Creation of alias from the Cloudera Manager UI fails on FIPS

7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

Users attempting to create aliases through the Cloudera Manager UI face issues in FIPS.

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

1. ssh to Knox host.
2. Export these directories: `export KNOX_GATEWAY_DATA_DIR="/var/lib/knox/gateway/data"; export KNOX_GATEWAY_CONF_DIR="/var/lib/knox/gateway/conf"`
3. Set the FIPS-specific options for the Knox CLI:

```
export KNOX_CLI_MEM_OPTS="--add-exports=java.base/sun.security.provider=bctls --add-exports=java.base/sun.security.provider=com.safelogic.cryptocomply.fips.core --add-modules=com.safelogic.cryptocomply.fips.core --add-modules=bctls --module-path=<BCTLS_JARS_DIR> -Dcom.safelogic.cryptocomply.fips.approved_only=true"
```

<BCTLS_JARS_DIR> is the directory containing the SafeLogic bctls and fips core jar files.

4. Run the following command to create the alias: `/opt/cloudera/parcels/CDH/lib/knox/bin/knoxccli.sh create-alias <ALIAS_NAME> <ALIAS_VALUE>`
5. Verify the addition using `/opt/cloudera/parcels/CDH/lib/knox/bin/knoxccli.sh list-alias`.

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

CDPD-71305: Concurrent impala shell connection failure

7.1.9 SP1 and its CHF, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

If a user makes a concurrent impala-shell connection through Knox, then the connection fails.

Use only one Knox role.

CDPD-73368: Knox token management is not working if Cookie Management is enabled

7.3.1, 7.3.1.100, 7.3.1.200

7.3.1.300

If Cookie Management is enabled, users are unable to access the Token Management page from the Knox Gateway UI by using KnoxSSO.

None.

Apache JIRA: [KNOX-3060](#)

CDPD-68146: Unable to update the log level for Knox from Cloudera Manager

7.1.9, 7.2.17, 7.2.18, 7.3.1, 7.3.1.100

7.3.1.200

Users are not able to change the log level for Knox from Cloudera Manager. Hence, it impacts debugging in case of any issue.

Change the level for the org.apache.knox.gateway logger in /var/lib/knox/gateway/conf/gateway-log4j2.xml file and restart Knox.

CDPD-60379: During rolling upgrade of Knox service, access fails with 503/500/404/403 error code

7.1.9, 7.2.18, 7.3.1

7.3.1

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

7.2.16, 7.2.17, 7.2.18, 7.1.7, 7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

At this time, Atlas does not communicate a log-out event with the external authentication management, Apache Knox. When you log out of Atlas, you can still open the instance of Atlas from the same web browser without re-authentication.

To prevent additional access to Atlas, close all browser windows and exit the browser.

CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes

7.1.7, 7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.1.7, 7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

None.

Knox issue with JDK version

jdk-1.8.0_391 is not supported.

Cloudera recommends using Cloudera supported JDKs.

CDPD-74843: Logs missing in third-party libraries

7.3.1

7.3.1.300

Some third-party libraries have missing logs due to a missing log4j library, which affects the ability to diagnose and troubleshoot issues. Knox is unable to modify the ROOT logger's level due to the missing log4j-slf4j-impl dependency.

Known Issues in Apache Kudu

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

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3:

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2:

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF 1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1

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

7.3.1.100

7.3.1

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](#).

Unsupported feature

Kudu HMS Sync is disabled and is not yet supported.

Known Issues in Livy

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

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified for Livy in this release.

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified for Livy in this release.

Known Issues identified in Cloudera Runtime 7.3.1

There are no new known issues identified for Livy in this release.

Known Issues in Navigator Encrypt

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

Known Issues in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1.300 SP1 CHF1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF 1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1

Failed to start navencrypt-mount.service

7.1.9, 7.3.1

After installing NavEncrypt on SLES, if the command `systemctl status navencrypt-mount` fails with the error: "Failed to start navencrypt-mount.service: Unit navencrypt-mount.service failed to load: No such file or directory", the systemd files need to be reloaded..

Run the command :

```
(sudo) systemctl daemon-reload
```

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

7.1.7, 7.1.9, 7.3.1

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

None

Known Issues in Apache Oozie

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

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF 1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1

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:

- 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.

Known Issues in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2

CDPD-87097: Installation fails with a NullPointerException (NPE) in the Ozone-Recon log

7.3.1.400, 7.3.1.500

This is a known intermittent issue in 7.3.1.400. NPE thrown does not tell the actual cause and this intermittent issue is not reproducible as this comes sometimes during first run of installation of the cluster. This looks to be more of an CDEP environment issue.

CDPD-87080: 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

7.3.1.400, 7.3.1.500

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"`.

For more information, see [Changing temporary path for Ozone services and CLI tools](#).

Known Issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF 1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1

CDPD-72557: Cannot use a single Ozone client to access encrypted data from multiple clusters

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

The Ozone client cannot be used to read or write encrypted data to multiple Ozone clusters using different KMS's for encrypted key management. This is because Ozone client uses KMS provider URI specified in the client side configurations instead of using the provider presented by multiple Ozone Managers.

None

CDPD-35141: Move operation is not supported across buckets in Ozone

Moving keys across buckets is currently not supported in Ozone.

Apache Issue: [HDDS-6585](#)

CDPD-54885: Ozone Prometheus does not work with TLS

The Prometheus service shipped by Ozone does not support TLS mode. So, Prometheus is not able to gather metrics from Ozone endpoints when TLS is enabled.

Go to `Ozone Configuration Ozone Prometheus Endpoint Token` and update the Ozone Prometheus Endpoint Token with any random string. This makes a token available in the process directory of the Ozone endpoint in plaintext, which Prometheus can use to get around the TLS limitation.

CDPD-69681: Ozone Ranger plugin does recursive listing for ACL checks

Ozone Ranger plugin does recursive listing in unoptimized path for performing ACL check on a directory. Recursive rename and delete operations on directories with layers of subdirectories causes a performance problem.

[Enable Ranger subaccess optimization for rename and delete operations in Ozone.](#)

Apache Issue: [HDDS-10823](#)

CDPD-51490: When you are upgrading from Cloudera Base on premises 7.1.8 or 7.1.9 to Cloudera Base on premises 7.3.1 and installing HDFS and Ozone datanode instances on the same server (in case of a small test or development platform), can cause the Ozone UI to fail as it attempts to use the same default port 9864 already taken by HDFS.

On the Ozone configuration, overwrite the `hdds.datanode.client.port` value = 19864 and restart the Ozone service. This will allow Ozone UI to load without colliding with the HDFS service.

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 Cloudera Base on premises 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 Cloudera Base on premises 7.1.9 SP1, there is a missing block location in the output of Ozone debug `chunkinfo` command for EC bucket.

None.

In Cloudera Base on premises 7.1.9 SP1 and Cloudera Base on premises 7.1.7 SP3, there is an intermittent key put failure after stopping the OM leader.

None.

In Cloudera Base on premises 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 Cloudera Base on premises to Cloudera Base on premises 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 Cloudera Base on premises 7.1.7 SP3 to Cloudera Base on premises 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 Cloudera Base on premises 7.1.7 SP3, complete Kerberos owner name with domain is used. However, in Cloudera Base on premises 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 OMs. 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 OMs to the certificate directories of all OMs.

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 Cloudera Base on premises 7.1.8 to Cloudera Base on premises 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 Cloudera Base on premises 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 Cloudera Base on premises 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.S3MultipartUploadCompleteReques
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteReques
at
org.apache.hadoop.ozone.protocolPB.OzoneManagerRequestHandler.handleWriteReques
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.runCommand(OzoneManag
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 Cloudera Base on premises 7.1.8 before upgrading to Cloudera Manager 7.11.3 with Cloudera Base on premises 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 Cloudera Base on premises 7.1.8 to Cloudera Base on premises 7.1.9.

OPSAPS-68159: If you did not deactivate the Ozone parcel 718.2.x on Cloudera Manager 7.7.1 with Cloudera Base on premises 7.1.8 before upgrading to Cloudera Manager 7.11.3 with Cloudera Base on premises 7.1.9, the Ozone roles during the Cloudera Base on premises 7.1.8 upgrade to Cloudera Base on premises 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 -f -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"`

For more information, see [Changing temporary path for Ozone services and CLI tools](#).

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.

For more information, see [Changing temporary path for Ozone services and CLI tools](#).

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 Cloudera 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 Cloudera 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" message 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 Cloudera Base on premises 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 temporary path for Ozone services and 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

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

Known issues identified in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.100 CHF1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1

There are no new known issues identified in this release.

Known Issues in Apache Phoenix

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

Known issues identified in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.100 CHF1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1

There are no new known issues identified in this release.

Known Issues in Apache Ranger

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

Known Issues in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1.300 SP1 CHF1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF 1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1

OPSAPS-75673: Wrong enablement of Ranger RMS Database Full Sync command

7.1.8, 7.1.9, 7.1.9 SP1, 7.2.18, 7.3.1

The Ranger RMS Database Full Sync command should be enabled only when all RMS server instances are stopped. This is required to ensure that the RMS database synchronizes correctly without introducing conflicts or data corruption. However, when HA (High Availability) is enabled on the cluster, the command becomes available from Cloudera Manager Ranger RMS Actions drop-down, even though only one Ranger RMS instance is stopped while the others are still running.

None.

UnsupportedClassVersionError

7.1.9, 7.1.9 SP1, 7.3.1

JDK 8 deployments support Nashorn JavaScript engine, which is built-in and fully compatible, whereas JDK 17 deployments support GraalJS script engine due to unavailability of Nashorn.

When your cluster supports both JDK 8 and JDK 17, then while a Java application, running on JDK 8, uses generic interfaces like ScriptEngine (from the javax.script package), the ScriptEngineManager class scans the classpath for available script engine implementations through the service provider mechanism, and detects the GraalJS as a provider for JavaScript, because in this case the GraalJS library (version 22.3.0) is also included on the classpath. The ScriptEngineManager then attempts to instantiate it, when requesting a "js" or "javascript" engine, and triggers an UnsupportedClassVersionError.

Remove the GraalJS library from the classpath.

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)

7.3.1

7.3.1.100

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

7.1.9 CHF3, 7.1.9 SP1, 7.3.1

7.1.9 SP1 CHF4, 7.3.1.100

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

7.1.9 SP1, 7.3.1

List command is listing all the tables even when the user permission is revoked. And also the command does not add any deny policy to Ranger for that specific user.

This behavior is currently not supported in HBase shell. Must be handled manually using the Ranger policy change.

CDPD-68739: The revoke command does not work when using the HBase shell**7.1.9 SP1, 7.3.1**

While using the HBase shell, running the revoke command does not cancel the user permission. Users are able to perform actions even after running the revoke command.

None.

CDPD-58704: hadoop roll key/key delete command shows operation failed error when one KMS host is down, even when operation succeeds**7.1.9, 7.1.9 SP1, 7.3.1**

In case of rollover/delete, client sends one more (last after delete request) request to KMS instances to clean their cache and that too to all registered kms instances. if one KMS instance is stopped (not deleted), the client gets a runtime exception.

This simply returns the runtime exception on client end for stopped instances but doesn't break any functionality.

CDPD-56803: When there is no existing policy for user and a revoke request comes from hbase, then will get this error**7.1.9, 7.1.9 SP1, 7.3.1**

None.

CDPD-56741: Improvement in log message when jwtauth not used**7.1.9, 7.1.9 SP1, 7.3.1**

None.

CDPD-56738: Ranger RMS showing FileNotFoundException: /usr/share/java/oraclepki.jar in Oracle 19 setup**7.1.9, 7.1.9 SP1, 7.3.1**

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**7.1.9****7.1.9 CHF2, 7.2.18**

If you were using multiple user search filters in Audit Access Tab on Ranger Admin UI, after upgrading to Cloudera Runtime 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**7.1.9, 7.1.9 SP1, 7.3.1**

KMS keys are not case sensitive.

No workaround. Such key combinations are very rare and the migration doc was updated to check such keys before starting the migration.

CDPD-42598: Kafka policy creation allowed with incorrect permissions

7.1.8, 7.1.9, 7.1.9 SP1, 7.3.1

When creating a Kafka policy from the UI, the permissions "Idempotent write" and "Cluster action" are not displayed as they are not applicable for the "topic" resource, but when creating a policy for the "topic" resource with the permissions "Idempotent write" and "Cluster Action", the policy is created successfully when the expected behaviour is that the policy creation must fail as the permission is not applicable for the Kafka topic resource.

None.

CDPD-41582: Atlas Resource Lookup : Classification for "entity-type" lists only classification for the following payload: {"resourceName": "classification", "userInput": "", "resources": {"classification": []}}

7.1.8, 7.1.9, 7.1.9 SP1, 7.2.16, 7.2.17, 7.2.18, 7.3.1

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

7.1.8, 7.1.9, 7.1.9 SP1, 7.3.1

7.3.1.200

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.

Known Issues in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1.300 SP1 CHF1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF 1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1

There are no new known issues identified in this release.

Known Issues in Schema Registry

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

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF 1

There are no new known issues identified in this release.

Cloudera Runtime 7.3.1

CDPD-40380: Authorization checking issue when Kerberos is disabled

7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

Due to an issue in Ranger, when Kerberos is disabled then it is not possible to check authorization.

1. Open Schema Registry configuration in Cloudera Manager.
2. Find the ranger.plugin.schema-registry.service.name field.
3. Replace GENERATED_RANGER_SERVICE_NAME with the actual name of the service.
4. Restart the Schema Registry service.

CDPD-49304: AvroConverter does not support composite default values

7.1.7, 7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

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

7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

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

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

7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

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

Known Issues in Apache Solr

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

Known issues identified in Cloudera Runtime 7.3.1.500 SP3

Ranger Audit Solr Data Lake restoration fails

7.3.1 and its higher versions

The following steps highlight a manual remediation process for a specific failure encountered during a Data Lake restore or upgrade, particularly affecting the `ranger_audit` Solr collection. This issue arises when Solr leaves behind residual replica instancedirs or data directory entries, causing recreation attempts to fail with an error stating "another core is already defined."



Note: Ensure that you use these steps whenever `restore-datalake` stalls on `rangerAuditsCollection` and you find the "another core is already defined there" error in Solr logs.

1. Check the Solr service state in Cloudera Manager and ensure that all Solr server roles display a green status.
2. Inspect the current `ranger_audits` state.

```
curl "<http://<gateway>>:8983/solr/admin/collections?
action=CLUSTERSTATUS&collection=ranger_audits&wt=json" | jq
'.cluster.collections.ranger_audits'
```

```
# Per Solr host
grep -A3 "Error CREATEing SolrCore" /var/log/solr/solr-cmf-*.log.out
# Look for `Could not create a new core ... as another core
is already defined there`.
```

3. Clean up any lingering shards or replicas.

- Remove collection state from Solr

```
curl "http://<gateway>:8983/solr/admin/collections?
action=DELETE&name=ranger_audits&wt=json"
```

If the collection is already absent, the command returns an error, which you can ignore.

- Check if there are any leftover `ranger_audits` directories on HDFS.

```
hdfs dfs -ls /solr-infra/ | grep ranger_audits
```

- Remove the `ranger_audits` data directories on HDFS.

```
hdfs dfs -rm -r /solr-infra/ranger_audits_shardX_replica_nY
```

- Check if there are any leftover `ranger_audits` directories.

```
ls -ld /var/lib/solr-infra/ranger_audits*
```

- Remove the local `ranger_audits` instance directories.

- Restart the Solr service.

Validate CLUSTERSTATUS again to ensure that ranger_audits is removed and the cluster is healthy before re-running the restore operation.

- Re-run the Data Lake restore operation.

Known issues identified in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.1.100 CHF1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1

HBase indexer does not load netty and snappy libraries

7.3.1 and its higher versions

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:

- Go to the Key-Value Store Indexer service Configuration .
- Search for the Key-Value Store Indexer Service Environment Advanced Configuration Snippet (Safety Valve) property.
- 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
```

- Restart the Key-Value Store Indexer service by clicking Key-Value Store Indexer service Actions Restart .

HBase Indexer does not work with JDK 17

7.3.1 and its higher versions

Depending on the Cloudera Manager version used with Cloudera, 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 jav
a.base/java.lang=ALL-UNNAMED --add-opens java.base/java.lang
.reflect=ALL-UNNAMED
```

Splitshard operation fails after CDH 6 to Cloudera upgrade

7.3.1 and its higher versions

Collections are not reindexed during an upgrade from CDH 6 to Cloudera 7 because Lucene 8 (Cloudera) 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(Solr
IndexSplitter.java:318) ~[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.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

7.3.1 and its higher versions

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.MasterRegistryFetchException: Exception making rpc to masters [quasar-bmyccr-2.quasar-bmyccr.root.hwx.site, 22001, -1]
    at org.apache.hadoop.hbase.client.MasterRegistry.lambda$groupCall$1(MasterRegistry.java:244)
    at org.apache.hadoop.hbase.util.FutureUtils.lambda$addListener$0(FutureUtils.java:68)
    at java.util.concurrent.CompletableFuture.uniWhenComplete(CompletableFuture.java:774)
    at java.util.concurrent.CompletableFuture.uniWhenCompleteStage(CompletableFuture.java:792)
    at java.util.concurrent.CompletableFuture.whenComplete(CompletableFuture.java:2153)
    at org.apache.hadoop.hbase.util.FutureUtils.addListener(FutureUtils.java:61)
    at org.apache.hadoop.hbase.client.MasterRegistry.groupCall(MasterRegistry.java:228)
    at org.apache.hadoop.hbase.client.MasterRegistry.call(MasterRegistry.java:265)
    at org.apache.hadoop.hbase.client.MasterRegistry.getMetaRegionLocations(MasterRegistry.java:282)
    at org.apache.hadoop.hbase.client.ConnectionImplementation.locateMeta(ConnectionImplementation.java:900)
    at org.apache.hadoop.hbase.client.ConnectionImplementation.locateRegion(ConnectionImplementation.java:867)
    at org.apache.hadoop.hbase.client.ConnectionImplementation.relocateRegion(ConnectionImplementation.java:850)
    at org.apache.hadoop.hbase.client.ConnectionImplementation.locateRegionInMeta(ConnectionImplementation.java:981)
    at org.apache.hadoop.hbase.client.ConnectionImplementation.locateRegion(ConnectionImplementation.java:870)
    at org.apache.hadoop.hbase.client.RpcRetryingCallerWithReadReplicas.getRegionLocations(RpcRetryingCallerWithReadReplicas.java:319)
    ... 21 more
Caused by: org.apache.hadoop.hbase.client.RetriesExhaustedException: Failed contacting masters after 1 attempts.
Exceptions:
java.io.IOException: Call to address=quasar-bmyccr-2.quasar-bmyccr.root.hwx.site/172.27.19.4:22001 failed on local exception: java.io.IOException: java.lang.RuntimeException: Found no valid authentication method from options
    at org.apache.hadoop.hbase.client.MasterRegistry.lambda$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.ZKConnectionRegistry'
```

Unable to see single valued and multivalued empty string values when querying collections after upgrade to Cloudera

7.3.1 and its higher versions

After upgrading from CDH or HDP to Cloudera, you are not able to see single valued and multi Valued empty string values in Cloudera.

This behavior in Cloudera 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

7.3.1 and its higher versions

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

7.3.1 and its higher versions

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

7.3.1 and its higher versions

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

7.3.1 and its higher versions

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

7.3.1 and its higher versions

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 Cloudera upgrade

7.3.1 and its higher versions

During an upgrade from CDH6 to Cloudera, collections are not reindexed because Lucene 8 (Cloudera) 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 Cloudera 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

7.3.1 and its higher versions

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:

- 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 Apache Spark, the impact or changes to the functionality, and the workaround.

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF1

The following section lists the known issues identified in this release:

CDPD-80239: Non-deterministic SQL expressions should set indeterminate map stage output level

7.3.1, 7.3.1.100 CFH1, 7.3.1.200 SP1, 7.3.1.300 SP1 CHF1, 7.3.1.400 SP2

Spark is supposed to handle non-deterministic keys, as long as they are marked with deterministic=false in their data type attributes. For Spark's random data this contract is not honored when there is a task failure. As a result, duplicate or missing data can be produced when the Spark executors are relaunched in new node managers.

Use the client configuration `spark.global.deterministic` to override any input-level deterministic configuration. If set to true, all inputs are deterministic, if set to false all inputs are indeterministic.

Known Issues identified in Cloudera Runtime 7.3.1

The following section lists the known issues identified in this release:

Spark 3: RAPIDS Accelerator is not available

7.3.1, 7.3.1.100 CHF1, 7.3.1.200 SP1, 7.3.1.300 SP1 CHF1, 7.3.1.400 SP2

The RAPIDS Accelerator for Apache Spark is currently not available in Cloudera Runtime 7.3.1. None.

The CHAR(n) type handled inconsistently, depending on whether the table is partitioned or not.

7.3.1

7.3.1.100 CHF1

In upstream Spark 3 the `spark.sql.legacy.charVarcharAsString` configuration was introduced, but it does not solve all incompatibilities with Spark 2.

None. A new configuration `spark.cloudera.legacy.charVarcharLegacyPadding` will be introduced in a future version to keep compatibility with Spark 2, but it isn't available in 7.3.1.



Note: The CHAR type is legacy in SQL, and using it is discouraged. Cloudera recommends using VARCHAR or STRING instead.

Apache Jira: [SPARK-33480](#)

Known Issues in Spark Atlas Connector

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

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1

There are no new known issues identified in this release.

Known Issues in Apache Sqoop

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

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3:

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2:

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1.300 SP1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF 1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1**CDPD-54770: Unable to read Sqoop metastore created by an older HSQLDB version**

7.1.9 SP1 and its CHF, 7.3.1 and its higher versions

If you have upgraded to Cloudera Base on premises 7.1.8 Cumulative hotfix 4 or higher versions, you may encounter issues in reading the Sqoop metastore that was created using an older version of HyperSQL Database (HSQLDB).

Cloudera upgraded the HSQLDB dependency from 1.8.0.10 to 2.7.1 and this causes incompatibility issues in Sqoop jobs that are stored in HSQLDB.

After upgrading to Cloudera Base on premises 7.1.8 Cumulative hotfix 4, you must upgrade the Sqoop metastore and convert the database files to a format that can easily be read by HSQLDB 2.7.1. For more information, see [Troubleshooting Apache Sqoop issues](#).

CDPD-44431: Using direct mode causes problems

7.1.9 SP1 and its CHF, 7.3.1 and its higher versions

Using direct mode has several drawbacks:

- Imports can cause an intermittent and overlapping input split.
- Imports can generate duplicate data.
- Many problems, such as intermittent failures, can occur.
- Additional configuration is required.

Stop using direct mode. Do not use the --direct option in Sqoop import or export commands.

Sqoop direct mode is disabled by default. However, if you still want to use it, enable it by either setting the `sqoop.enable.deprecated.direct` property globally in Cloudera Manager for Sqoop or by specifying it in the command-line through `-Dsqoop.enable.deprecated.direct=true`.

CDPD-3089: Avro, S3, and HCat do not work together properly

7.1.7 SP1 and its CHF, 7.1.9 SP1 and its CHF

Importing an Avro file into S3 with HCat fails with Delegation Token not available.

Parquet columns inadvertently renamed

Problem: Column names that start with a number are renamed when you use the `--as-parquetfile` option to import data.

Prepend column names in Parquet tables with one or more letters or underscore characters.

Importing Parquet files might cause out-of-memory (OOM) errors

Problem: Importing multiple megabytes per row before initial-page-run check (ColumnWriter) can cause OOM. Also, rows that vary significantly by size so that the next-page-size check is based on small rows, and is set very high, followed by many large rows can also cause OOM.

PARQUET-99

Known issues in Streams Messaging Manager

Learn about the known issues in Streams Messaging Manager, the impact or changes to the functionality, and the workaround.

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF 1

There are no new known issues identified in this release.

Cloudera Runtime 7.3.1**CDPD-74620: Data Explorer returns 500 Server Error when switching partitions**

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

In the Streams Messaging Manager Data Explorer, switching from a partition with a higher end offset to a partition with a lower end offset causes a 500 Server Error. This happens because the interface incorrectly applies the end offset of the previous partition to the new partition. As a result, the user interface displays No Data or an error message indicating that the start offset must be smaller than the log end offset.

Manually set the Offset in Data Explorer to a valid value for the partition you are viewing.

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

7.1.7, 7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

View the Streams Messaging Manager UI logs on the host.

CDPD-39313: Some numbers are not rendered properly in Streams Messaging Manager UI

7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

Very large numbers can be imprecisely represented on the UI. For example, bytes larger than 8 petabytes would lose precision.

None.

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

7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

Streams Messaging Manager cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You need to override bootstrap server URL (hostname:port as set in the listeners for broker). Add the bootstrap server details in Streams Messaging Manager safety valve in the following path:

1. In Cloudera Manager, go to SMMConfigurationStreams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml.
2. Add the following value for bootstrap servers.

```
streams.messaging.manager.kafka.bootstrap.servers=<comma-separated list of brokers>
```

3. Save your changes.
4. Restart Streams Messaging Manager.

CDPD-78694: Streams Messaging Manager UI cannot show replication details

7.3.1

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

The Kafka cluster replication details view is not available on the Cluster Replications page because the UI does not pass the required duration URL parameter when calling the Streams Messaging Manager REST Admin server.

You can query replication details by manually calling the Streams Messaging Manager REST Admin server with the required duration URL parameter.

```
api/v2/admin/replication-stats?duration=LAST_ONE_HOUR
```

CDPD-82560: Streams Messaging Manager known unprotected endpoint

7.1.7, 7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

The Streams Messaging Manager Frontend exposes the /cm-configs and /configs public endpoints without requiring authentication. These endpoints do not share any sensitive information.

None.

CDPD-88488: The Streams Messaging Manager server fails to start if its keystore password starts with a special character or includes quotation marks

7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

Ensure that the Streams Messaging Manager server keystore password meets the following requirements:

- It does not start with a special character. Special characters are permitted in other positions.
- It does not contain quotation marks (single or double) anywhere.

CDPD-76375: Topic column name is incorrect in the replications table

7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

On the **Cluster Replications** page, in the replications table, the topic column name is incorrectly named **Source Topic Name** instead of Target Topic Name. This means that the **Source Topic Name** column currently shows the topic name on the target cluster, instead of the topic name on the source cluster.

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, the impact or changes to the functionality, and the workaround.

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF 1

There are no new known issues identified in this release.

Cloudera Runtime 7.3.1

CDPD-22089: Streams Replication Manager does not sync re-created source topics until the offsets have caught up with target topic

7.1.7, 7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

Messages written to topics that were deleted and re-created are not replicated until the source topic reaches the same offset as the target topic. For example, if at the time of deletion and re-creation there are a 100 messages on the source and target clusters, new messages will only get replicated once the re-created source topic has 100 messages. This leads to messages being lost.

None

CDPD-11079: Blacklisted topics appear in the list of replicated topics

7.1.7, 7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

If a topic was originally replicated but was later disallowed (blacklisted), it will still appear as a replicated topic under the /remote-topics REST API endpoint. As a result, if a call is made to this endpoint, the disallowed topic will be included in the response. Additionally, the disallowed topic will also be visible in the Streams Messaging Manager UI. However, its Partitions and Consumer Groups will be 0, its Throughput, Replication Latency and Checkpoint Latency will show N/A.

None

CDPD-30275: Streams Replication Manager may automatically re-create deleted topics on target clusters

7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

If `auto.create.topics.enable` is enabled, deleted topics might get automatically re-created on target clusters. This is a timing issue. It only occurs if remote topics are deleted while the replication of the topic is still ongoing.

1. Remove the topic from the topic allowlist with `srm-control`. For example:

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGET_CLUSTER] --remove [TOPIC1]
```

2. Wait until Streams Replication Manager is no longer replicating the topic.
3. Delete the remote topic in the target cluster.

CDPD-80872: Streams Replication Manager replication-records-lag is incorrect on empty partitions with non-zero end offset

7.1.9, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

When a replicated partition is empty and its end offset is non-zero, the `replication-records-lag` metric is incorrectly reported as the end offset (instead of 0 or NaN).

Wait for a new message to be written into the partition. When it gets replicated, the metric is reported correctly.

Limitations

Streams Replication Manager cannot replicate Ranger authorization policies to or from Kafka clusters

Due to a limitation in the Kafka-Ranger plugin, Streams Replication Manager cannot replicate Ranger policies to or from clusters that are configured to use Ranger for authorization. If you are using Streams Replication Manager to replicate data to or from a cluster that uses Ranger, disable authorization policy synchronization in Streams Replication Manager. This can be achieved by clearing the Sync Topic Acls Enabled (`sync.topic.acls.enabled`) checkbox.

Streams Replication Manager cannot ensure the exactly-once semantics of transactional source topics

Streams Replication Manager 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 Streams Replication Manager 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 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
```

Streams Replication Manager checkpointing is not supported for transactional source topics

Streams Replication Manager 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 and Yarn Queue Manager

Learn about the known issues in Yarn and Yarn Queue Manager, the impact or changes to the functionality, and the workaround.

Known Issues identified in Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Known Issues identified in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1.200 SP1

The following section lists the known issues identified in this release:

COMPX-14682: Fix health check after Queue Manager restart

7.1.9, 7.1.9 SP1, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

In some cases, after the QM webapp and config-service are restarted Cloudera Manager reports a healthy status after some time. However, the configuration service might not be ready and usable yet when this status is reported. Therefore, QM (webapp) is not ready and usable yet.

Wait for a minute or two after restart and then try QM.

Known Issues identified in Cloudera Runtime 7.3.1.100 CHF 1

There are no new known issues identified in this release.

Known Issues in Cloudera Runtime 7.3.1**COMPX-14820: Delete Queue and its Children throws "Queue capacity was reduced to zero, but failed to delete queue."**

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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'

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

Products:

- Cloudera Manager for Cloudera on premises

- Cloud Base Cloudera Manager for Cloudera on 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

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

When migrating from an H2 database to a PostgreSQL database in YARN Queue Manager after installation or upgrade, you might encounter an issue only when you have followed the following specific scenario:

- New install or upgrade to Cloudera 7.1.9, forcing migration from H2 to PostgreSQL database.
- Upgrade to Cloudera 7.1.9 CHF2, moving back to H2 database.
- Upgrade to Cloudera 7.1.9 SP1 with valid PostgreSQL connection details in Queue Manager configurations.

To avoid any issues during the upgrade to version Cloudera 7.1.9 SP1, ensure that PostgreSQL connection details are removed from the YARN database configuration if you prefer to continue using the H2 database.

CDPD-56559: MapReduce jobs can intermittently fail during a rolling upgrade.

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

During a rolling upgrade between Cloudera 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 Cloudera 7.1.9 to a higher version.

None.

COMPX-12021 Queue Manager configurations on Scheduler Configuration page are not working

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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.

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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.

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

When 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 Clusters is deleted or upgraded or restored.

COMPX-6628: Unable to delete single leaf queue assigned to a partition.

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

In the current implementation, if there is an existing managed queue in Relative mode, then conversion to Weight mode is not be 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

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

MapReduce application framework is loaded from HDFS instead of being present on the NodeManagers. By default the `mapreduce.application.framework.path` property is set to the appropriate value, but third-party applications with their own configurations does not launch.

Set the `mapreduce.application.framework.path` property to the appropriate configuration for third-party applications.

JobHistory URL mismatch after server relocation

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

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

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

When the Job ACL feature is enabled using Cloudera Manager (YARN Configuration Enabl 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.

YARN Resource Manager UI does not display logs

7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500

When the parameter mapreduce.cluster.acl.enabled is set to true, the Yarn RM UI does display and logs and the Logs are not available message is displayed.

Set mapreduce.cluster.acl.enabled to false.

Unsupported Features

The following YARN features are currently not supported in Cloudera:

- Application Timeline Server v2 (ATSv2)
- Auxiliary Services
- Container Resizing
- Distributed or Centralized Allocation of Opportunistic Containers
- Distributed Scheduling
- Docker on YARN (DockerContainerExecutor)
- 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.

Cloudera Runtime 7.3.1.500 SP3

There are no new known issues identified in this release.

Cloudera Runtime 7.3.1.400 SP2

There are no new known issues identified in this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no new known issues identified in this release.

Cloudera Runtime 7.3.1.200 SP1

There are no new known issues identified in this release.

Cloudera Runtime 7.3.1.200CHF 1

There are no new known issues identified in this release.

Cloudera Runtime 7.3.1

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.

Cloudera Runtime 7.3.1.500 SP3

Summary:

A new option to ignore spark_process.attributes, details and sparkPlanDescription is introduced.

Previous behavior:

The spark_process entity attributes details and sparkPlanDescription are populated with query plan details, which can contain a large amount of text, often in megabytes. This amount of data can incur unnecessary processing costs.

New behavior:

The attribute atlas.notification.consumer.preprocess.spark_process.attributes is set to true by default to avoid populating the details and sparkplandescription attributes. These attributes can be ignored using above configurations in Atlas server. Ignoring these attributes helps to eliminate the cost of having large amount of data processed in Atlas.

To keep the data for details and sparkplandescription, atlas.notification.consumer.preprocess.spark_process.attributes, set them to false explicitly.

Cloudera Runtime 7.3.1.400 SP2

Summary:

A new option to ignore spark_process.attributes, details and sparkPlanDescription is introduced.

Previous behavior:

The spark_process entity attributes details and sparkPlanDescription are populated with query plan details, which can contain a large amount of text, often in megabytes. This amount of data can incur unnecessary processing costs.

New behavior:

The attribute atlas.notification.consumer.preprocess.spark_process.attributes is set to false by default. Set it to true to avoid populating the details and sparkplandescription attributes. These attributes can be ignored using above configurations in Atlas server. Ignoring these attributes helps to eliminate the cost of having large amount of data processed in Atlas.

Summary:

A new option to avoid sending details and sparkPlanDescription in the Spark process entity is introduced.

Previous behavior:

The spark_process entity attributes details and sparkPlanDescription are populated with query plan details, which can contain a large amount of text, often in megabytes. This amount of data can incur unnecessary processing costs.

New behavior:

The atlas.spark.plan.enabled is set to true by default. Set it to false to send the details and sparkPlanDescription attributes in the Spark process entity. When these attributes are not sent, the cost of having large amount of data processed in Atlas is avoided.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.1**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 Apache Avro

Behavioral changes denote a marked change in behavior from the previously released version to this version of Apache Avro.

Cloudera Runtime 7.3.1.500 SP3

There are no behavioral changes this release.

Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes this release.

Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes this release.

Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes this release.

Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes this release.

Cloudera Runtime 7.3.1

There are no behavioral changes this release.

Behavioral Changes in Cloud Connectors

Behavioral changes denote a marked change in behavior from the previously released version to this version of Cloud Connectors.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

There are no behavioral changes in this release.

Behavioral Changes in Cruise Control

Behavioral changes denote a marked change in behavior from the previously released version to this version of Cruise Control.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

There are no behavioral changes in this release.

Behavioral Changes in HBase

Behavioral changes denote a marked change in behavior from the previously released version to this version of HBase.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

Add security headers to Thrift/HTTP server

Summary:

Added security headers to the Thrift or HTTP server.

Previous behavior:

Previously, security headers were absent in the responses of Thrift or HTTP servers when utilizing HBase-exposed services.

New behavior:

To address this, additional security headers are introduced in the responses of the HBase Thrift server when HTTP or HTTPS transport is enabled.

Apache JIRA: [HBASE-27118](#)

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

There are no behavioral changes in this release.

Behavioral Changes in HDFS

Behavioral changes denote a marked change in behavior from the previously released version to this version of HDFS.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

There are no behavioral changes in this release.

Behavioral Changes in Hive

Behavioral changes denote a marked change in behavior from the previously released version to this version of Apache Hive.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

Summary:

Hive Metastore now uses dynamic leader election and automated housekeeping.

Previous behavior:

The Hive Metastore (HMS) previously used a static, host-based method for leader election, as the `metastore.housekeeping.leader.election` property defaulted to `host`. Additionally, critical housekeeping and compaction processes were turned off by default.

New behavior:

The HMS now uses a dynamic, lock-based method for leader election, as the `metastore.housekeeping.leader.election` property now defaults to `lock`. In addition, housekeeping and compaction processes are now enabled by default, ensuring that the Metastore automatically manages these critical operations in a single node across the warehouse.

Summary:

Increased batch sizes for COMPUTE STATS

Previous behavior:

The COMPUTE STATS query previously failed on tables containing more than 5000 columns. This issue was specific to wide tables and could not be resolved by dropping and rerunning the query.

New behavior:

To resolve this, we enable the batch retrieval or insertion of the object metadata by default value of the `hive.metastore.direct.sql.batch.size` property is changed from 0 to 1000, and the default value of the `metastore.rawstore.batch.size` property is changed from -1 to 500. After this change, COMPUTE STATS queries now run successfully on tables with more than 5000 columns.

Summary:

Removal of the `engine.hive.enabled` property

Previous behavior:

The engine.hive.enabled property was set to "true" to enable the creation of Hive-enabled Iceberg table.

New behavior:

The engine.hive.enabled property is removed because Hive is already supported and SQL engines are not required to explicitly specify this property.

Summary:

Change in the way dates are parsed from string by ignoring trailing invalid characters

Previous behavior:

Prior to this release, SQL functions or date operations involving invalid dates returned "null".

New behavior:

Now, a valid date is extracted and returned from a string value if there is a valid date prefix in the string. This change partially restores the behavior introduced in HIVE-20007 and makes the handling of trailing invalid characters more consistent.

The following table illustrates the behavior changes before and after the fix:

Strong value	Behavior (before HIVE-20007)	Previous behavior (after HIVE-20007)	Current behavior (after HIVE-27586)
2023-08-03_16:02:00	2023-08-03	null	2023-08-03
2023-08-03-16:02:00	2023-08-03	null	2023-08-03
2023-08-0316:02:00	2024-06-11	null	2023-08-03
03-08-2023	0009-02-12	null	0003-08-20
2023-08-03 GARBAGE	2023-08-03	2023-08-03	2023-08-03
2023-08-03TGARBAGE	2023-08-03	2023-08-03	2023-08-03
2023-08-03_GARBAGE	2023-08-03	null	2023-08-03

This change affects various Hive SQL functions and operators that accept dates from string values, such as CAST (V AS DATE), CAST (V AS TIMESTAMP), TO_DATE, DATE_ADD, DATE_DIFF, WEEKOFYEAR, DAYOFWEEK, and TRUNC.

Summary:

Change in the way date and timestamp values are parsed.

Previous behavior:

Some Hive date and timestamp functions, such as unix_timestamp(), from_unixtime(), date_format(), and cast() used the SimpleDateFormat class for printing and parsing date and timestamp objects.

New behavior:

A new configurable hive.datetime.formatter property is introduced that enables you to choose between SimpleDateFormat and DateTimeFormatter for the unix_timestamp, from_unixtime, and date_format SQL functions. The default value is set to 'DATETIME'.

Summary:

Change in value for the hive.map.groupby.sorted property

Previous behavior:

The value of the hive.map.groupby.sorted property was set to 'true'.

New behavior:

The value of the `hive.map.groupby.sorted` property is changed to 'false' to disable optimization. The change was introduced to address data correctness issues noticed in query results on tables with `CLUSTER BY` and `SORT BY`.

Apache Jira: [HIVE-27876](#)

Summary:

Disabling join disjunctive predicate pushdown

Previous behavior:

With `hive.optimize.join.disjunctive.transitive.predicates.pushdown` enabled by default, queries with disjunctive predicates could cause HiveServer2 to crash or run out of memory during compilation.

Apache Jira: [HIVE-28310](#)

New behavior:

The `hive.optimize.join.disjunctive.transitive.predicates.pushdown` setting is now disabled by default, enhancing HiveServer2 stability and preventing crashes and out-of-memory errors. In some rare cases, queries with joins and unions become slightly less efficient but the difference should not be noticeable by the end-users.

Summary:

Hive CBO fallback strategy configuration

Previous behavior:

The `hive.cbo.fallback.strategy` property was set to `CONSERVATIVE` by default. In case of an error during the cost-based optimizer phase, Hive would fallback to the legacy optimizer, potentially reducing optimization efficiency and masking serious or unrecoverable errors.

Apache Jira: [HIVE-27831](#)

New behavior:

The default value for `hive.cbo.fallback.strategy` is now set to `NEVER`. Hive no longer falls back to the legacy optimizer and cost-based optimizer errors are fatal. Hidden compilation errors will now show up immediately and additional actions are required to compile and execute the query successfully.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

Summary:

Timestamp Conversion Handling

Previous behavior:

In Cloudera Base on premises 7.1.9 version, `hive.strict.timestamp.conversion` was set to `true`, enforcing strict failures when casting between timestamp and numeric types.

Apache Jira

Starting from Cloudera Base on premises 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'.

Apache Jira

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.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

Summary:

Simplified Impala INSERT performance analysis

Previous behavior:

Query profiles for INSERT operations displayed a long and misleading duration for 'Got Metastore client' right before 'Fired Metastore events'.

New behavior:

Now, the catalog timeline accurately includes a dedicated item for 'Prepared InsertEvent data' within the 'Catalog Server Operation' section of Impala query profiles. This allows users to clearly see the time spent collecting file checksums during INSERT operations, providing more precise performance insights.

Apache Jira: [IMPALA-13960](#)

Summary:

Optimized ephemeral storage for Impala's Catalogd

Previous behavior:

The ephemeral storage limit of 512MB often caused container evictions due to insufficient space, especially during JVM heap dumps.

New behavior:

The ephemeral storage limit for Catalogd is now set dynamically based on the JVM heap size (Xmx), calculated as

$$2 * Xmx + \text{Buffer space}$$

This approach ensures adequate storage for heap dumps and minimizes the risk of container eviction, resulting in improved stability for Catalogd and other services sharing node resources.

Summary:

Cleanup subdirectories in `truncate/insert overwrite` if recursing listing is enabled

Previous behavior:

Impala did not consistently delete files located in subdirectories of external tables during `TRUNCATE` and `INSERT OVERWRITE` operations, even when recursive listing was enabled. This led to leftover data in subdirectories after these operations, resulting in data corruption.

New behavior:

Now, directories are also deleted in addition to (non-hidden) data files, with the exception of hidden and ignored directories. Also, setting `DELETE_STATS_IN_TRUNCATE=false` is no longer supported by default when truncating non-transactional tables; attempting this will result in an exception. If the old behavior is absolutely required, you can set the `--truncate_external_tables_with_hms` flag to false, but be aware that this will also reintroduce the bug that was fixed by this change.

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

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1**Summary:**

Impala Query Analysis Behavior with Ranger.

Previous behavior:

Impala previously verified `WRITE` access for the service user on HDFS table/partition(s) during query analysis of `INSERT` and `LOAD DATA` statements in legacy catalog mode. Permissions were computed based on HDFS settings, including ACLs, when tables and partitions were instantiated.

New behavior:

To address performance concerns, HDFS permissions are now skipped during query analysis. The service user is assumed to have `READ_WRITE` access to all HDFS paths associated with the target table when Ranger is enabled. Ranger policies remain enforced during query execution for `INSERT` and `LOAD DATA` statements, ensuring security compliance.

Apache Jira: [IMPALA-11871](#)

Summary:

Expression rewrite behavior for Hive views with auto-generated column aliases.

Previous behavior:

Impala attempted to simplify `CAST` expressions for all columns, including those with Hive auto-generated aliases (such as `_c0`), introduced by the `SimplifyCastExprRule` optimization in [IMPALA-10836](#). In views created in Hive without explicit column aliases, this could lead to `AnalysisException` errors during query execution. For example, a view using `CAST` on a column labeled as `_c0` might fail with:

```
AnalysisException: Could not resolve column/field reference:
'failing_view._c0'
```

New behavior:

Impala now skips rewriting expressions that are associated with Hive auto-generated column aliases (for example, `_c0`, `_c1`, etc.). This preserves the correct column mapping across nested views and

avoids errors during query analysis. This change allows queries to succeed without requiring you to explicitly rename columns in Hive views.

Apache Jira: [IMPALA-11871](#)

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

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 Hue

Behavioral changes denote a marked change in behavior from the previously released version to this version of Hue.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3:

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2:

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

There are no behavioral changes in this release.

Behavioral Changes in Iceberg

Behavioral changes denote a marked change in behavior from the previously released version to this version of Apache Iceberg.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

There are no behavioral changes in this release.

Behavioral Changes in Kafka

Behavioral changes denote a marked change in behavior from the previously released version to this version of Kafka.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

The `LdapLoginModule` is blocked by default

Summary:

The `LdapLoginModule` (`com.sun.security.auth.module.LdapLoginModule`) class is blocked by default and can not be used in SASL JAAS configurations.

Previous behavior:

The `LdapLoginModule` class was an allowed login module.

New behavior:

The `LdapLoginModule` class is no longer allowed and is blocked by default. Use the `PlainLoginModule` (`org.apache.kafka.common.security.plain.PlainLoginModule`) class for LDAP authentication.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

There are no behavioral changes in this release.

Behavioral Changes in Knox

Behavioral changes denote a marked change in behavior from the previously released version to this version of Apache Knox.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

Knox token impersonation config

Summary

Knox token service has been changed to use the identity assertion provider configuration for impersonation.

Previous behavior

The token service had its own impersonation configuration.

New behavior

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 behavior

The name of the file was `gateway-identity.pem`.

New behavior

The name of the file is now `gateway-client-trust.pem`.

Composite authorization provider misconfiguration

Summary

Composite authorization provider misconfiguration behavior

Previous behavior

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 behavior

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 behavior

Requests for topologies which are not yet fully deployed result in HTTP 404 responses.

New behavior

Requests for topologies which are not yet fully deployed result in HTTP 503 responses.

Behavioral Changes in Kudu

Behavioral changes denote a marked change in behavior from the previously released version to this version of Kudu.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3:

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2:

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

There are no behavioral changes in this release.

Behavioral Changes in Livy

Behavioral changes denote a marked change in behavior from the previously released version to this version of Livy.

Behavioral changes in Livy in Cloudera Runtime 7.3.1.500 SP3

Summary:

Removed the CustomAuthenticationProvider class.

Previous behavior:

The CustomAuthenticationProvider class was available.

New behavior:

The CustomAuthenticationProvider class is no longer available. A client program may be interrupted by a NoSuchMethodError exception.

Summary:

Removed the NoneAuthenticationProvider class.

Previous behavior:

The NoneAuthenticationProvider class was available.

New behavior:

The NoneAuthenticationProvider class is no longer available. A client program may be interrupted by a NoSuchMethodError exception.

Behavioral changes in Livy in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral changes in Livy in Cloudera Runtime 7.3.1.300 SP1 CHF1

There are no behavioral changes in this release.

Behavioral changes in Livy in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral changes in Livy in Cloudera Runtime 7.3.1.100 CHF1

There are no behavioral changes in this release.

Behavioral changes in Livy in Cloudera Runtime 7.3.1

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 on premises version 7.3.1.

Behavioral Changes in Oozie

Behavioral changes denote a marked change in behavior from the previously released version to this version of Oozie.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

There are no behavioral changes in this release.

Behavioral Changes in Ozone

Behavioral changes denote a marked change in behavior from the previously released version to this version of Ozone.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

There are no behavioral changes in this release.

Behavioral Changes in Apache Parquet

Behavioral changes denote a marked change in behavior from the previously released version to this version of Apache Parquet.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

There are no behavioral changes in this release.

Behavioral Changes in Phoenix

Behavioral changes denote a marked change in behavior from the previously released version to this version of Phoenix.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

There are no behavioral changes in this release.

Behavioral Changes in Ranger

Behavioral changes denote a marked change in behavior from the previously released version to this version of Apache Ranger.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

Summary: Hive authorization from Ranger for Alter Table Rename command does not require CREATE database permission on the database where the renamed table will be created.

Previous behavior:

In releases earlier than 7.3.1.100, whenever Alter Table Rename command was used across databases in Hive, authorization from Ranger required CREATE database permission for the user on the target database in which the renamed table was created.

New behavior:

In 7.3.1 CHF1 and later releases, whenever Alter Table Rename command is used across databases in Hive, authorization from Ranger does not check for CREATE database permission for the user on the target database in which the renamed table will be created.

Behavioral Changes in Cloudera Runtime 7.3.1

Summary: Policy resources in the Ranger Admin UI are being added using React JS instead of Backbone JS

Previous behavior:

Earlier with Backbone JS, when you copied and pasted resource values containing commas or spaces (for example, in Hive policy resources: database1, database2), the UI automatically split them into separate values — database1 and database2. The same behaviour applied to space-separated values. Because of this, you were not allowed to enter resource names containing commas and spaces, and this limitation affected all service policy resources.

New behavior:

After upgrading from Backbone JS to React JS, this restriction has been removed. Now, React JS treats pasted values with commas or spaces as a single entry. Hence, you can no longer paste multiple values at once; you must manually add each resource value.

Summary: 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 Schema Registry

Behavioral changes denote a marked change in behavior from the previously released version to this version of Schema Registry.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

There are no behavioral changes in this release.

Behavioral Changes in Solr

Behavioral changes denote a marked change in behavior from the previously released version to this version of Solr.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

Summary:

The default value of the `hbaseindexer.httpserver.disabled` environment parameter changed from false to true.

Previous behavior:

You needed to change the value of the `hbaseindexer.httpserver.disabled` environment parameter to true to switch off the REST interface. This was necessary to prevent use of the `--http` argument when using the `hbase-indexer` command line tool. Using the `--http` argument for the `hbase-indexer` command line tool to invoke Lily indexer through REST API allowed adding/listing/removing indexers with any user without the need for authentication.

New behavior:

The HBase Lily indexer REST API is switched off by default.

Behavioral Changes in Sqoop

Behavioral changes denote a marked change in behavior from the previously released version to this version of Sqoop.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3:

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2:

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

There are no behavioral changes in this release.

Behavioral Changes in Streams Messaging Manager

Behavioral changes denote a marked change in behavior from the previously released version to this version of Streams Messaging Manager.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

There are no behavioral changes in this release.

Behavioral Changes in Streams Replication Manager

Behavioral changes denote a marked change in behavior from the previously released version to this version of Streams Replication Manager.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

There are no behavioral changes in this release.

Behavioral Changes in Spark

Behavioral changes denote a marked change in behavior from the previously released version to this version of Spark.

Behavioral changes in Spark in Cloudera Runtime 7.3.1.500 SP3

Summary:

Added super-interface `java.nio.file.PathMatcher` to the `IOFileFilter` API.

Previous behavior:

None.

New behavior:

Unless updated to use the new super-interface, recompilation of a client program may be terminated with the following message: A client class C is not abstract and does not override abstract method in `java.nio.file.PathMatcher`.

Summary:

Added `java.io.IOException` exception thrown to the `CountingInputStream` class in the `CountingInputStream.afterRead` API.

Previous behavior:

None.

New behavior:

Unless updated to handle the new exception, recompilation of a client program may be terminated with the following message: Unreported exception `java.io.IOException` must be caught or declared to be thrown.

Summary:

Added `java.io.FileNotFoundException` exception thrown to the `FileUtils.isFileOlder` class in the `FileUtils` API.

Previous behavior:

None.

New behavior:

Unless updated to handle the new exception, recompilation of a client program may be terminated with the following message: Unreported exception `java.io.FileNotFoundException` must be caught or declared to be thrown.

Summary:

Return value type has been changed in the `StreamIterator` class in the `StreamIterator.iterator` API.

Previous behavior:

Return value was `java.util.Iterator<T>`.

New behavior:

New return value is `org.apache.commons.io.StreamIterator<T>`

Summary:

Removed the `java.io.EOFException` exception thrown in the `SwappedDataInputStream` class in the `SwappedDataInputStream.skipBytes` API.

Previous behavior:

Application code could handle the `java.io.EOFException` exception thrown.

New behavior:

Unless removed from the application code, recompilation of a client program may be terminated with the following message: `Cannot override skipBytes() in org.apache.commons.io.input.SwappedDataInputStream; overridden method does not throw java.io.EOFException.`

Summary:

Return value type has been changed in the `StreamIterator` class in the `StreamIterator.iterator` API.

Previous behavior:

Return value was `java.util.Iterator<T>`

New behavior:

New return value is `org.apache.commons.io.StreamIterator<T>`.

The previous method has been removed because the return type is part of the method signature. A client program may be interrupted by a `NoSuchMethodError` exception.

Behavioral changes in Spark in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral changes in Spark in Cloudera Runtime 7.3.1.300 SP1 CHF1

There are no behavioral changes in this release.

Behavioral changes in Spark in Cloudera Runtime 7.3.1.200 SP1**Summary:**

Rebase Spark3 to Apache Spark 3.5.4 in Cloudera Runtime.

Previous behavior:

Spark 3.4.1 was the default version in Cloudera Runtime.

New behavior:

Spark 3.5.4 is the default Spark version in Cloudera Runtime.



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

Behavioral changes in Spark in Cloudera Runtime 7.3.1.100 CHF1

There are no behavioral changes in this release.

Behavioral changes in Spark in Cloudera Runtime 7.3.1

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](#)

Behavioral Changes in Spark Atlas Connector

Behavioral changes denote a marked change in behavior from the previously released version to this version of Spark Atlas Connector.

Behavioral changes in Cloudera Runtime 7.3.1.500 SP3

There are no behavioral changes in Spark Atlas Connector in this release.

Behavioral changes in Cloudera Runtime 7.3.1.400 SP2

Summary: Spark plan is optional in Spark process entity sent to Atlas

Previous behavior:

Atlas would experience an Out-of-Memory error during processing if the spark_process entity's details and sparkPlanDescription fields contained too large strings.

New behavior:

A new Cloudera Manager property atlas.spark.plan.enabled was added to enable users to enable or disable the details and sparkPlanDescription field. (Default value: true.)

When set to true (default), the fields can be used as before. When set to false, it reduces memory usage, but removes the fields from the Atlas UI.

Behavioral changes in Cloudera Runtime 7.3.1

There are no behavioral changes in Spark Atlas Connector in this release.

Behavioral Changes in Yarn and Yarn Queue Manager

Behavioral changes denote a marked change in behavior from the previously released version to this version of Yarn and Yarn Queue Manager.

Behavioral Changes in Cloudera Runtime 7.3.1.500 SP3

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.400 SP2

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.300 SP1 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.200 SP1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1.100 CHF 1

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Runtime 7.3.1

There are no behavioral changes in this release.

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 Cloudera release. Marking an item as deprecated gives you time to plan for removal in a future Cloudera release.

Moving

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

Removed

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

Removed Components and Product Capabilities

- Apache Spark 2

Spark 3 is the default Spark version in Cloudera Runtime. Spark 2 (and Livy 2) has been removed and no longer available in 7.3.1

**Important:**

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

Refer to *Upgrading Spark* for more information on upgrading Spark clusters to 7.3.1.0, and *Migrating Spark Applications* for more information on migrating your existing Spark applications between versions 2 and 3.

- Apache Livy 2 (see [Deprecation Notices for Apache Livy](#))
- Apache Zeppelin (see [Deprecation Notices for Apache Zeppelin](#))

Please contact Cloudera Support or your Cloudera Account Team if you have any questions.

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:

- PostgreSQL 12
- MariaDB 10.4
- MySQL 5.7

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:

- RHEL 8.6
- RHEL 7.9
- RHEL 7.9 (FIPS)
- CentOS 7.9
- SLES 12 SP5
- 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 Cloudera. For a full list of deprecation and/or removals in the version Apache Kafka shipped with Cloudera Runtime, review the *Notable Changes* as well as the *Release Notes* on <https://kafka.apache.org/>.

Deprecated

MirrorMaker (MM1)

MirrorMaker is deprecated. Cloudera recommends that you use Streams Replication Manager instead.

--zookeeper

The --zookeeper option is only supported for the kafka-configs tool and should be only used when updating SCRAM Credential configurations. The --zookeeper option is either deprecated in or removed from other Kafka command line tools. Cloudera recommends that you use the --bootstrap-server option instead.

Deprecation Notices for 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 Cloudera Base on premises 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 Cloudera Base on premises 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 Cloudera 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 on premises.

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.

Technical Service Bulletins

Technical Service Bulletins and Customer Advisories for the Cloudera 7.3.1 release and its Service Packs and Cumulative Hotfixes.

TSB 2025-820: Potential Data Integrity Issues Found in Ozone

Learn more about the details communicated in TSB-820.

Summary

The Cloudera Engineering team has identified the following data integrity issues with Apache Ozone (Ozone):

1. In certain situations, handling of failure paths when recovering from disk hardware failures, disk full situations, or over-replication can result in the incorrect deletion of some storage containers on those disk(s). In rare cases, all replicas of the container can be affected, leading to the data within that container becoming unavailable. Under certain extreme conditions, permanent data loss could occur.

Reference: CDPD-83416

2. A bug in the snapshot deep cleaning service and the object deletion path can lead to potential missing blocks of a snapshot key. This can happen only for the keys that were deleted from the active object store after the snapshot was created.

Reference: CDPD-83417

Component(s) affected

- Ozone

Releases affected

- Cloudera 7.3.1

Addressed in release/refresh/patch

- Cloudera 7.3.1.300 SP1 CHF 1

Knowledge Base article

For the latest update on this issue see the corresponding Knowledge Base article: [TSB 2025-820: Potential Data Integrity Issues Found in Ozone](#)

TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager

Learn more about the details communicated in TSB-835.

Summary

Executing the "Dry Run" action for Ozone replication schedules with a "Listing type" of "Incremental only" or "Incremental with fallback to full file listing" will result in a run where the changes are not replicated and also omitted from the subsequent replication runs.

Unless a "Full file listing" replication run is executed, the changes made between the dry run and the previous run are not replicated to the target. Such a scenario may occur when, during the dry run action of an Ozone replication policy with INCREMENTAL_ONLY and INCREMENTAL_WITH_FALLBACK_TO_FULL_FILE_LISTING replication type, generates a temporary snapshot on the source, which doesn't get deleted. On the next incremental run, all changes that occurred on the source Ozone bucket between the last successful run and the last dry

run operation, will go unnoticed by the Replication Manager. This situation results in the failure to replicate such changes, to the destination Ozone bucket.

Component(s) affected

- Cloudera Replication Manager

Releases affected

- Cloudera Platform 7.3.1
- Cloudera Base on premises 7.1.9

Addressed in release/refresh/patch

- Cloudera Manager for Cloudera Platform 7.3.1
 - Cloudera Manager 7.13.1.400 (Dry Run feature is temporarily disabled)
- Cloudera Manager for Private Cloud Data Services
 - Cloudera Manager 7.11.3 CHF16 (Dry Run feature is temporarily disabled)

Knowledge Base article

For the latest update on this issue see the corresponding Knowledge Base article: [TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager](#)

Apache Parquet CVE-2025-30065

Learn more about the details communicated in TSB-847.

Background

On April 1, 2025, a critical vulnerability in the parquet-avro module of Apache Parquet ([CVE-2025-30065](#), [CVSS score 10.0](#)) was announced.

Cloudera has determined the list of affected products, and is issuing this TSB to provide details of remediation for affected versions.

Upgraded versions are being released for all currently affected [supported releases](#) of Cloudera products. Customers using older versions are advised to upgrade to a [supported release](#) that has the remediation, once it becomes available.

Vulnerability Details

Exploiting this vulnerability is only possible by modifying the accepted schema used for translating Parquet files and subsequently submitting a specifically crafted malicious file.

[CVE-2025-30065](#) | Schema parsing in the parquet-avro module of Apache Parquet 1.15.0 and previous versions allows bad actors to execute arbitrary code.

Impact

Schema parsing in the parquet-avro module of Apache Parquet 1.15.0 and previous versions allows bad actors to execute arbitrary code. Attackers may be able to modify unexpected objects or data that was assumed to be safe from modification. Deserialized data or code could be modified without using the provided accessor functions, or unexpected functions could be invoked.

Deserialization vulnerabilities most commonly lead to undefined behavior, such as memory modification or remote code execution.

Releases Affected

- Cloudera Runtime 7.3.1.100 CHF1 and lower versions

Resolved In

- Cloudera Runtime 7.3.1.200 SP1

Mitigation

Until Cloudera has released product version with the Apache Parquet vulnerability fix, please continue to use the the mitigations listed below:

Customers with their own FIM Solution:

1. Utilize a File Integrity Monitoring (FIM) solution. This allows administrators to monitor files at the filesystem level and receive alerts on any unexpected or suspicious activity in the schema configuration.

General advisory:

1. Use network segmentation and traffic monitoring with a device capable of deep packet inspection, such as a network firewall or web application firewall, to inspect all traffic sent to the affected endpoints. Configure alerts for any suspicious or unexpected activity. You may also configure sample analysis parameters to include:
 - a. Parquet file format “magic bytes” = PAR1
 - b. Connections from sending hosts that are not expected source IP ranges.
2. Be cautious with Parquet files from unknown or untrusted sources. If possible, do not process files with uncertain origins or that can be ingested from outside the organization.
3. Ensure that only authorized users have access to endpoints that ingest Parquet files.

Knowledge Base article

For the latest update on this issue see the corresponding Knowledge Base article: [TSB 2025-847: Apache Parquet CVE-2025-30065](#)

Fixed Common Vulnerabilities and Exposures

Common vulnerabilities and Exposures (CVEs) fixed in Cloudera Runtime 7.3.1 release, and its Service Packs and Cumulative hotfixes.

Cloudera Runtime 7.3.1.500 SP3

Common vulnerabilities and Exposures (CVEs) fixed in Cloudera Runtime 7.3.1.500 SP3.

- [CVE-2020-13956](#) - HttpClient
- [CVE-2022-1471](#) - Snakeyaml
- [CVE-2023-39410](#) - Apache Avro
- [CVE-2023-46589](#) - Apache Tomcat
- [CVE-2024-13009](#) - Eclipse Jetty
- [CVE-2024-21634](#) - AWS Java SDK Bundle
- [CVE-2024-23672](#) - Apache Tomcat
- [CVE-2024-24549](#) - Apache Tomcat
- [CVE-2024-34750](#) - Apache Tomcat
- [CVE-2024-8184](#) - Eclipse Jetty
- [CVE-2024-47561](#) - Apache Avro
- [CVE-2024-53990](#) - Async HTTP Client

Cloudera Runtime 7.3.1.400 SP2

Common vulnerabilities and Exposures (CVEs) fixed in Cloudera Runtime 7.3.1.400 SP2.

- [CVE-2024-36124](#) - iq80.snappy

- [CVE-2025-27553](#) - commons-vfs2

Cloudera Runtime 7.3.1.300 SP1 CHF 1

Common vulnerabilities and Exposures (CVEs) fixed in Cloudera Runtime 7.3.1.300 SP1 CHF 1.

There are no CVEs fixed in this release.

Cloudera Runtime 7.3.1.200 SP1

Common vulnerabilities and Exposures (CVEs) fixed in Cloudera Runtime 7.3.1.200 SP1.

- [CVE-2024-8391](#) - Vertx
- [CVE-2025-30065](#) - Apache Parquet

Cloudera Runtime 7.3.1.100 CHF 1

Common vulnerabilities and Exposures (CVEs) fixed in Cloudera Runtime 7.3.1.100 CHF 1.

- [CVE-2021-47621](#) Classgraph
- [CVE-2021-41973](#) Mina-core
- [CVE-2023-22102](#) mysql-connector-j

Cloudera Runtime 7.3.1

Common vulnerabilities and Exposures (CVEs) fixed in Cloudera Runtime 7.3.1.

- [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